

NEWTON'S METHOD AND SYMMETRY FOR SEMILINEAR ELLIPTIC PDE ON THE CUBE

JOHN M. NEUBERGER, NÁNDOR SIEBEN, AND JAMES W. SWIFT

ABSTRACT. We seek discrete approximations to solutions $u : \Omega \rightarrow \mathbb{R}$ of semilinear elliptic partial differential equations of the form $\Delta u + f_s(u) = 0$, where f_s is a one-parameter family of nonlinear functions and Ω is a domain in \mathbb{R}^d . The main achievement of this paper is the approximation of solutions to the PDE on the cube $\Omega = (0, \pi)^3 \subseteq \mathbb{R}^3$. There are 323 possible isotropy subgroups of functions on the cube, which fall into 99 conjugacy classes. The bifurcations with symmetry in this problem are quite interesting, including many with 3-dimensional critical eigenspaces. Our automated symmetry analysis is necessary with so many isotropy subgroups and bifurcations among them, and it allows our code to follow one branch in each equivalence class that is created at a bifurcation point. Our most complicated result is the complete analysis of a degenerate bifurcation with a 6-dimensional critical eigenspace.

This article extends the authors' work in *Automated Bifurcation Analysis for Nonlinear Elliptic Partial Difference Equations on Graphs* (Int. J. of Bifurcation and Chaos, 2009), wherein they combined symmetry analysis with modified implementations of the gradient Newton-Galerkin algorithm (GNGA, Neuberger and Swift) to automatically generate bifurcation diagrams and solution graphics for small, discrete problems with large symmetry groups. The code described in the current paper is efficiently implemented in parallel, allowing us to investigate a relatively fine-mesh discretization of the cube. We use the methodology and corresponding library presented in our paper *An MPI Implementation of a Self-Submitting Parallel Job Queue* (Int. J. of Parallel Prog., 2012).

1. INTRODUCTION

We are interested in finding and approximating solutions $u : \Omega \rightarrow \mathbb{R}$ of semilinear elliptic equations with zero Dirichlet boundary conditions,

$$(1) \quad \begin{cases} \Delta u + f_s(u) = 0 & \text{in } \Omega \\ u = 0 & \text{on } \partial\Omega, \end{cases}$$

where $f_s : \mathbb{R} \rightarrow \mathbb{R}$ satisfies $f_s(0) = 0$, $f'_s(0) = s$, and Ω is a region in \mathbb{R}^2 or \mathbb{R}^3 . Our code also works for zero Neumann boundary conditions, and a wide range of nonlinearities. In this paper we present results for PDE (1) with $f_s(t) = st + t^3$ on the square $\Omega = (0, \pi)^2$, and more challengingly, on the cube $\Omega = (0, \pi)^3$. By finding and following new, bifurcating branches of (generally) lesser symmetry we are able to approximate, within reason, any solution that is connected by branches to the trivial branch. The more complicated solutions bifurcating farther from the origin $(u, s) = (0, 0)$ are of course progressively more challenging to locate and accurately approximate.

Generally, we apply Newton's method to the gradient of an action functional whose critical points are solutions to our PDE. For an exposition of our initial development of the gradient Newton-Galerkin algorithm (GNGA) and our first application of it to the square, see [16].

This article extends the methods for so-called partial difference equations (PdE) from [13] to large graphs, that is, fine mesh discretizations for PDE. For small graphs with possibly large symmetry groups, the code in [13] automated the analysis of symmetry, isotypic decomposition, and bifurcation. We use here the GAP (Groups, Algorithms, and Programming, see [8]) and Mathematica codes presented in those articles to automatically generate a wealth of symmetry

2000 *Mathematics Subject Classification.* 20C35, 35P10, 65N25.

Key words and phrases. Bifurcation with Symmetry, Semilinear Elliptic PDE, Cube, GNGA, Numerical. January 31, 2013.

information for use by our branch-following C++ code. Some of this information is summarized in the bifurcation digraph, which shows the generic symmetry-breaking bifurcations. In [13] we developed two modified implementations of the gradient Newton-Galerkin algorithm, namely the tangent algorithm (tGNGA) for following bifurcation curves and the cylinder algorithm (cGNGA) for switching branches at bifurcation points. Together with a secant method for locating these bifurcation points, in the current PDE setting we are able to handle most difficulties that arise when encountering accidental degeneracies and high-dimensional critical eigenspaces.

Since we use here a fine mesh to investigate PDE (1) on the cube $\Omega = (0, \pi)^3$, the practical implementation of the algorithms from [13] requires increased efficiencies. In particular, we use isotypic decompositions of invariant fixed-point spaces to take advantage of the block diagonal structure of the Hessian matrix. In doing so, we substantially reduce the dimension of the Newton search direction linear system, and hence also reduce the number of costly numerical integrations. The same theory allows for reduced dimensions in many of the search spaces when seeking new, bifurcating solutions near high-dimensional bifurcation points in the presence of symmetry.

Even with these efficiency improvements, we found it necessary to convert our high-level branch following strategy to use parallel computing. Most of the details of the parallel implementation can be found in [15], where we present a general methodology using self-submitting job queues to implement many types of mathematical algorithms in parallel. In particular, therein we develop a light-weight, easy-to-use C++ parallel job queue library, which we have used in obtaining the cube results found in this article.

Our numerical results are summarized in bifurcation diagrams, which plot a scalar function such as the value of u at a generic point $u(x^*, y^*, z^*)$ versus s , for approximate solutions to Equation (1) with parameter s . These diagrams indicate by line type the Morse Index (MI) of solutions, which typically changes at bifurcation and turning points. We present graphics for individual approximate solutions in several formats. For the most part, we find that representative graphics using a small, fixed collection of patches (“flags”) most clearly show the symmetries of real-valued functions of three variables. We call these *flag diagrams*. We also include some contour plots of actual solution approximations. A fairly comprehensive collection of graphics and supporting information describing the symmetries of functions on the cube, possible bifurcations of nonlinear PDE on the cube, and more example approximate PDE solutions can be found on the companion website [14].

In [13], we considered many small graphs where scaling was not used, and hence PDE were not involved. In the present setting, we approximate a solution u to PDE (1) with a vector $\mathbf{u} = (u_n) \in \mathbb{R}^N$ whose components represent u values at N regularly spaced grid points in Ω located a distance Δx apart. Thus, our approach is equivalent to applying our algorithms to the finite dimensional semilinear elliptic partial difference equation (PdE)

$$(2) \quad -L\mathbf{u} + f_s(\mathbf{u}) = 0 \quad \text{in } \Omega_N,$$

where Ω_N is a graph with N vertices coming from a grid. The matrix L is in fact the graph Laplacian on Ω_N , scaled by $\frac{1}{(\Delta x)^2}$, and modified at boundary vertices to enforce a zero-Dirichlet problem boundary condition. See [11] for a discussion of ghost points for enforcing boundary conditions. For general regions in \mathbb{R}^2 and \mathbb{R}^3 , we approximate eigenvectors of L using standard linear techniques, e.g., Matlab’s *eigs* or some other easy to use implementation of ARPACK. For the square and cube the eigenfunctions are of course well known explicitly in terms of sine functions, so the consideration of L is not necessary. Since accurate PDE results require the dimension N to be very large, in the expansions of our approximate solutions $\mathbf{u} = \sum_{m=1}^M a_m \psi_m$ we use $M \ll N$ discretized eigenfunctions of $-\Delta$ with this boundary condition.

In Section 2 we present some theory for the action functional, its gradient and Hessian, symmetry of functions, the corresponding fixed-point subspaces, and bifurcations with symmetry. We apply the general symmetry theory to the basis generation process for the square and cube. In Section 3 we outline the algorithms used in our project. We include a high-level description of our numerical methods and corresponding implementations, including the new use of self-submitting parallel job

queues applied to obtain accurate high-resolution solutions for the cube. We also describe our method for taking advantage of the block structure of the Hessian and our procedure for generating contour plots of approximate solutions. Section 4 contains the results from our experiment on the square, essentially an efficient and automatic refinement of the computations found in [16].

Our main numerical results are found in Section 5. Namely, we investigate PDE (1) on the cube, where it is required to use a large number of eigenmodes and spacial grid points in order to find nodally complicated solutions of high MI, lying in many different fixed-point spaces of the fairly large symmetry group. We present several interesting examples from the companion website. The website shows examples of a solution with each of the symmetry types that we found in our investigation. Due to space limitations, we do not show all of these solutions in this paper. Rather, we concentrate on the first six primary bifurcations, and one of the secondary bifurcations with T_d symmetry (the symmetry of a tetrahedron). Two of the primary bifurcations that we consider have degenerate bifurcation points, including one with a six-dimensional eigenspace that is the direct sum of two 3-dimensional irreducible representations of \mathbb{O}_h , the symmetry of the cube.

There does not seem to be much in the literature that specifically investigates the bifurcation and symmetry of solutions to semilinear elliptic PDE on the cube. The article [3] is interesting for pushing the nonlinearity power to the critical exponent in the cube case. The interested reader can consult works by Zhou and co-authors for alternate but related methods and algorithms for computing solutions to semilinear elliptic PDE, e.g., [18, 19, 20] and the recent book [5].

2. SYMMETRY AND INVARIANCE

For the convenience of the reader, in this section we summarize enough notation and theory from [13] to follow our new results. We also include square and cube-specific information required to apply our algorithms in our particular cases.

2.1. The Functional Setting. Our techniques rely on two levels of approximation, namely the restriction of functions to a suitably large M -dimensional Galerkin subspace $B_M \subseteq H = H_0^1(\Omega)$, and the discretization of Ω to Ω_N . We call the natural numbers M and N the Galerkin and spacial dimensions of our approximations, respectively. For the regions Ω considered in this paper, it suffices to divide the region up into N squares or cubes with edge length Δx , and then place a gridpoint x_n in the center of each such cell. The graph Ω_N has a vertex v_i corresponding to each gridpoint, with edges e_{nj} determined by the several neighbors x_j which are at distance Δx away from x_n . With this arrangement, the simple numerical integration scheme used below in Equations 6 and 8 to evaluate the nonlinear terms in our gradient and Hessian computations becomes the midpoint method.

The eigenvalues of the negative Laplacian with the zero Dirichlet boundary condition satisfy

$$(3) \quad 0 < \lambda_1 < \lambda_2 \leq \lambda_3 \leq \dots \rightarrow \infty,$$

and the corresponding eigenfunctions $\{\psi_m\}_{m \in \mathbb{N}}$ can be chosen to be an orthogonal basis for the Sobolev space H , and an orthonormal basis for the larger Hilbert space $L^2 = L^2(\Omega)$, with the usual inner products. In the cases $\Omega = (0, \pi)^d$ for $d = 2$ and $d = 3$, we take the first M eigenvalues (counting multiplicity) from

$$\{\lambda_{i,j} := i^2 + j^2 \mid i, j \in \mathbb{N}\} \text{ or } \{\lambda_{i,j,k} := i^2 + j^2 + k^2 \mid i, j, k \in \mathbb{N}\},$$

and singly index them in a vector $\lambda = (\lambda_1, \dots, \lambda_M)$. In these cases, the corresponding eigenfunctions ψ_m that we use are appropriate linear combinations of the well-known eigenfunctions $\psi_{i,j}(x, y) = \frac{2}{\pi} \sin(ix) \sin(jy)$ and $\psi_{i,j,k}(x, y, z) = \left(\frac{2}{\pi}\right)^{3/2} \sin(ix) \sin(jy) \sin(kz)$. We process the eigenfunctions using the projections given in Section 2.3 in order to understand and exploit the symmetry of functions in terms of the nonzero coefficients of their eigenfunction expansions. The ψ_m are discretized as $\psi_m \in \mathbb{R}^N$, $m \in \{1, \dots, M\}$, by evaluating the functions at the gridpoints, i.e., $\psi_m = (\psi_m(x_1), \dots, \psi_m(x_N))$. The eigenvectors ψ_m form an orthonormal basis for an M -dimensional subspace of \mathbb{R}^N .

Using variational theory, we define a nonlinear functional J whose critical points are the solutions of PDE (1). We use the Gradient-Newton-Galerkin-Algorithm (GNGA, see [16]) to approximate these critical points, that is, we seek approximate solutions u lying in the subspace

$$B_M := \text{span}\{\psi_1, \dots, \psi_M\} \subseteq H,$$

which in turn are discretely approximated in \mathbb{R}^N by

$$\mathbf{u} = \sum_{m=1}^M a_m \boldsymbol{\psi}_m.$$

The coefficient vectors a in \mathbb{R}^M (simultaneously the approximation vectors \mathbf{u} in \mathbb{R}^N) are computed by applying Newton's method to the eigenvector expansion coefficients of the approximation $-L + f_s(\mathbf{u})$ of the gradient $\nabla J_s(u)$.

Let $F_s(p) = \int_0^p f_s(t) dt$ for all $p \in \mathbb{R}$ define the primitive of f_s . We then define the action functional $J_s : H \rightarrow \mathbb{R}$ by

$$(4) \quad J_s(u) = \int_{\Omega} \frac{1}{2} |\nabla u|^2 - F_s(u) dV = \sum_{m=1}^M \frac{1}{2} a_m^2 \lambda_m - \int_{\Omega} F_s(u) dV.$$

The class of nonlinearities f_s found in [1, 4] imply that J_s is well defined and of class C^2 on H . Computing directional derivatives and integrating by parts gives

$$(5) \quad J'_s(u)(\psi_m) = \int_{\Omega} \nabla u \cdot \nabla \psi_m - f_s(u) \psi_m dV = a_m \lambda_m - \int_{\Omega} f_s(u) \psi_m dV,$$

for $m \in \{1, \dots, M\}$. Replacing the nonlinear integral term with a sum that is in fact the midpoint method given our specific (square or cube) grid gives the gradient coefficient vector $g \in \mathbb{R}^M$ defined by

$$(6) \quad g_m = a_m \lambda_m - \sum_{n=1}^N f_s(u_n)(\boldsymbol{\psi}_m)_n \Delta V.$$

Here, the constant mesh area or volume factor is given by $\Delta V = \text{Vol}(\Omega)/N = \pi^d/N$. The functions $P_{B_M} \nabla J_s(u)$ and $\sum_{m=1}^M g_m \boldsymbol{\psi}_m$ are approximately equal and are pointwise approximated by the vector $-L\mathbf{u} + f_s(\mathbf{u})$.

To apply Newton's method to find a zero of g as a function of a , we compute the coefficient matrix h for the Hessian as well. A calculation shows that

$$(7) \quad J''_s(u)(\psi_l, \psi_m) = \int_{\Omega} \nabla \psi_l \cdot \nabla \psi_m - f'_s(u) \psi_l \psi_m dV = \lambda_l \delta_{lm} - \int_{\Omega} f'_s(u) \psi_l \psi_m dV,$$

where δ_{lm} is the Kronecker delta function. Again using numerical integration, for $l, m \in \{1, \dots, M\}$ we compute elements of h by

$$(8) \quad h_{lm} = \lambda_l \delta_{lm} - \sum_{n=1}^N f_s(u_n)(\boldsymbol{\psi}_l)_n (\boldsymbol{\psi}_m)_n \Delta V.$$

The coefficient vector $g \in \mathbb{R}^M$ and the $M \times M$ coefficient matrix h represent suitable projections of the L^2 gradient and Hessian of J , restricted to the subspace B_M , where all such quantities are defined. The least squares solution χ to the M -dimensional linear system $h\chi = g$ always exists and is identified with the projection of the search direction $(D_2^2 J_s(u))^{-1} \nabla_2 J_s(u)$ onto B_M . The L_2 search direction is not only defined for all points $u \in B_M$ such that the Hessian is invertible, but is in that case equal to $(D_H^2 J_s(u))^{-1} \nabla_H J_s(u)$.

The Hessian function $h_s : \mathbb{R}^M \rightarrow \mathbb{R}^M$ or $h : \mathbb{R}^{M+1} \rightarrow \mathbb{R}^M$ is very important for identifying bifurcation points. If $h(a^*, s^*)$ is invertible at a solution (a^*, s^*) , then the Implicit Function Theorem guarantees that there is locally a unique solution branch through (a^*, s^*) . When $h(a^*, s^*)$ is not

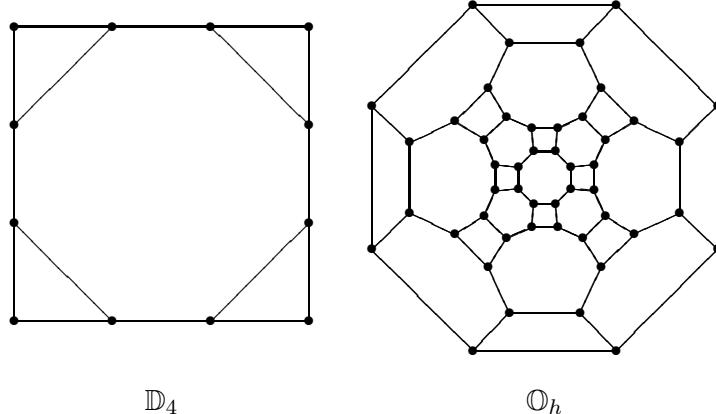


FIGURE 1. Small graphs used to generate symmetry information for analyzing functions on the square and cube, respectively. The graph on the right, with full octahedral symmetry, is the skeleton graph of the great rhombicuboctahedron, which has 48 vertices, 72 edges, and 26 faces. This solid can be inscribed in the cube, with 8 vertices in each of the 6 faces. Thus, there is a one-to-one correspondence between elements of \mathbb{O}_h and vertices of the skeleton graph, after one vertex (chosen arbitrarily) has been assigned to the identity element.

invertible, then (a^*, s^*) is a candidate for a bifurcation point, defined in the next subsection. The kernel

$$\tilde{E} = \text{Null } h(a^*, s^*)$$

of the Hessian at a bifurcation point is called the *critical eigenspace*. A Lyapunov-Schmidt reduction of the gradient $g(a^*, s^*)$ can be done to obtain the *Lyapunov-Schmidt reduced gradient* $\tilde{g} : \tilde{E} \rightarrow \tilde{E}$ [9]. Local to the point where h is singular, there is a one-to-one correspondence between zeros of g and zeros of \tilde{g} . The *Lyapunov-Schmidt reduced bifurcation equations* are $\tilde{g} = 0$. We refer to the reduced gradient or reduced bifurcation equations when the Lyapunov-Schmidt reduction is understood.

Newton's method in coefficient space is implemented by fixing s , initializing the coefficient vector a with a guess, and iterating

$$(9) \quad a \leftarrow a - \chi, \quad \text{where } h(a, s)\chi = g(a, s).$$

When it converges, the algorithm converges to vectors a and $\mathbf{u} = \sum_{m=1}^M a_m \psi_m$ giving $g = 0$, and hence an approximate solution to PDE (1) has been found. The search direction χ in Newton's method is found by solving the system in (9) without inverting $h(a, s)$. The solver we use returns the least squares solution for an overdetermined system and the solution with smallest norm for an underdetermined system. We observe experimentally that Newton algorithms work well even near bifurcation points where the Hessian is not invertible. In Section 3 we include brief descriptions of the tGNGA and cGNGA, the modifications of the GNGA actually implemented in our current code.

In Sections 2.3 and 2.4 we explain the details of our method for constructing a specific basis of eigenfunctions that allows our code to exploit and report symmetry information. Regardless of whether we know a basis for B_M in terms of sines or must approximate one using numerically computed eigenvectors of the sparse matrix L , we are required to first compute various symmetry quantities relevant to the region Ω and the possible symmetry types of expected solutions. For this we use GAP [8]. We start with a graph G which has the same symmetry as Ω_N , but with a significantly smaller number of vertices. In Figure 1 we show a 12-vertex graph with the \mathbb{D}_4 symmetry of the square, and a 48-vertex graph with the \mathbb{O}_h symmetry of the cube. These graphs are the smallest we found that allow functions on the vertices to have all possible symmetry types.

For example, the square with four vertices does not allow a function with only four-fold rotational symmetry. The group information used to analyze the symmetry of all functions on the square and cube regions was done automatically by our suite of programs. The automated GAP processes adapted from [13] would fail if fed instead the exceedingly large graphs Ω_N . Our GAP code is applied to the smaller graph G to generate all the files encoding the bifurcation digraph and the underlying fixed-point space decompositions. The results from the small-graph analysis inform us and our code about the symmetries of functions on Ω and the vectors in \mathbb{R}^N which approximate such functions at the gridpoints.

2.2. Symmetry of functions. We assume that f_s is odd; the case where f_s is not odd is easily inferred. To discuss the symmetry of solutions to Equation (1), we note that $\text{Aut}(\Omega) \times \mathbb{Z}_2 \cong \text{Aut}(\Omega_N) \times \mathbb{Z}_2 \cong \text{Aut}(G) \times \mathbb{Z}_2$, where G is one of the small graphs depicted in Figure 1. We define

$$\Gamma_0 = \text{Aut}(\Omega_N) \times \mathbb{Z}_2,$$

where $\mathbb{Z}_2 = \{1, -1\}$ is written multiplicatively.

The natural action of Γ_0 on \mathbb{R}^N is defined by

$$(10) \quad (\gamma \cdot \mathbf{u})_i = \beta u_{\pi^{-1}(i)},$$

where $\gamma = (\alpha_\pi, \beta) \in \Gamma_0$ and $\mathbf{u} \in \mathbb{R}^N$. We usually write α for $(\alpha, 1)$ and $-\alpha$ for $(\alpha, -1)$. The *symmetry* of \mathbf{u} is the isotropy subgroup $\text{Sym}(\mathbf{u}) := \text{Stab}(\mathbf{u}, \Gamma_0) = \{\gamma \in \Gamma_0 \mid \gamma \cdot \mathbf{u} = \mathbf{u}\}$. The symmetries $\text{Sym}(u)$ of functions $u : \Omega \rightarrow \mathbb{R}$ are isomorphic to $\text{Sym}(\mathbf{u})$. Two subgroups Γ_i and Γ_j of Γ_0 are called *conjugate* if $\Gamma_i = \gamma \Gamma_j \gamma^{-1}$ for some $\gamma \in \Gamma_0$. The *symmetry type* of \mathbf{u} is the conjugacy class $[\text{Sym}(\mathbf{u})]$ of the symmetry of \mathbf{u} ; a similar definition holds for the symmetry type of a function $u : \Omega \rightarrow \mathbb{R}$. We say that two symmetry types are *isomorphic* if they have isomorphic representatives. We use the notation $\mathcal{G} := \{\Gamma_0, \dots, \Gamma_q\}$ for the set of symmetries and $\mathcal{S} := \{S_0 = [\Gamma_0], \dots, S_r\}$ for the set of symmetry types.

Let X be the set of all solutions (\mathbf{u}, s) to (2) in $\mathbb{R}^N \times \mathbb{R}$. We define a *branch of solutions* to be a maximal subset of X that is a C^1 manifold with constant symmetry. The *trivial branch* $\{(0, s) \mid s \in \mathbb{R}\}$ contains the *trivial solution* $\mathbf{u} = 0$, which has symmetry Γ_0 if f_s is odd, and symmetry $\text{Aut}(G)$ otherwise. A *bifurcation point* is a solution in the closure of at least two different solution branches. We call the branch containing the bifurcation point the *mother*, and the other branches, for which the bifurcation point is a limit point, are called *daughters*. Note that there is not a bifurcation at a fold point, where a branch of constant symmetry is not monotonic in s .

The action of Γ_0 on \mathbb{R}^N induces an action of Γ_0 on \mathbb{R}^M , given the correspondence of functions and coefficient vectors $\mathbf{u} = \sum_{m=1}^M \psi_m$. The gradient function $g_s : \mathbb{R}^M \rightarrow \mathbb{R}^M$ is Γ_0 -equivariant, i.e., $g_s(\gamma \cdot a) = \gamma \cdot g_s(a)$ for all $\gamma \in \Gamma_0$, $a \in \mathbb{R}^M$, and $s \in \mathbb{R}$. As a consequence, if (a, s) is a solution to (2) then $(\gamma \cdot a, s)$ is also a solution to (2), for all $\gamma \in \Gamma_0$. Following the standard treatment [9, 13], for each $\Gamma_i \leq \Gamma_0$ we define the *fixed-point subspace* of the Γ_0 action on \mathbb{R}^M to be

$$\text{Fix}(\Gamma_i, \mathbb{R}^M) = \{a \in \mathbb{R}^M \mid \gamma \cdot a = a \text{ for all } \gamma \in \Gamma_i\}.$$

The fixed-point subspaces for any of the function spaces $V = G_M$, \mathbb{R}^n , \tilde{E} , or H is defined as

$$\text{Fix}(\Gamma_i, V) = \{u \in V \mid \gamma \cdot a = a \text{ for all } \gamma \in \Gamma_i\}.$$

There is a one-to-one correspondence between $a \in \mathbb{R}^M$ and $u \in G_M$, and we will often write $\text{Fix}(\Gamma_i)$ when the equation is valid for any ambient space. These fixed-point subspaces are important because they are g_s -invariant, meaning that $g_s(\text{Fix}(\Gamma_i)) \subseteq \text{Fix}(\Gamma_i)$. For efficiency in our code, we restrict g_s to one of these fixed-point subspaces, as described in Subsection 3.2.

As mentioned in the previous subsection, the reduced bifurcation equations are $\tilde{g} = 0$, where $\tilde{g} : \tilde{E} \rightarrow \tilde{E}$ is the reduced gradient on the critical eigenspace \tilde{E} . A fold point (a^*, s^*) is not a bifurcation point even though $h(a^*, s^*)$ is singular.

Consider a bifurcation point (a^*, s^*) in coefficient space, or the corresponding (\mathbf{u}^*, s^*) , and let $\Gamma_i = \text{Sym}(u^*)$ be the symmetry of the mother solution. Then there is a natural action of Γ_i on

\tilde{E} , and the reduced gradient \tilde{g} is equivariant. That is, $\tilde{g}(\gamma \cdot e, s) = \gamma \cdot \tilde{g}(e, s)$ for all $\gamma \in \Gamma_i$, and $e \in \tilde{E}$. Let Γ'_i be the kernel of the action of Γ_i on \tilde{E} . Then Γ_i/Γ'_i acts freely on \tilde{E} , and we say that the mother branch *undergoes a bifurcation with Γ_i/Γ'_i symmetry*. This bifurcation is *generic*, or *non-degenerate*, if the action of Γ_i/Γ'_i on \tilde{E} is irreducible, and other non-degeneracy conditions (see [9]) are met. For each $\Gamma_j \leq \Gamma_i$, if $e \in \text{Fix}(\Gamma_j, \tilde{E})$, then $\text{Sym}(\mathbf{u}^* + e) = \Gamma_j$. Most of the fixed-point subspaces in \tilde{E} are empty. The subgroups $\{\Gamma_j \mid \text{Fix}(\Gamma_j, \tilde{E}) \neq \emptyset\}$ can be arranged in a *lattice of isotropy subgroups*. The Equivariant Branching Lemma (EBL), described in [9], states that there is generically an *EBL branch* of bifurcating solutions with symmetry Γ_j if the fixed-point subspace $\text{Fix}(\Gamma_j, \tilde{E})$ is one-dimensional. In a gradient system such as PDE (1), there is generically a branch of bifurcating solutions with symmetry Γ_i if Γ_i is a maximal isotropy subgroup. See [9, 13] for details.

The *bifurcation digraph*, defined in [13], summarizes some information about all of the generic bifurcations that are possible for a system with a given symmetry. In particular, if there is a daughter with symmetry Γ_j created at a generic bifurcation of a mother solution with symmetry Γ_i in a gradient system, then there is an arrow in the bifurcation digraph

$$[\Gamma_i] \xrightarrow{\Gamma_i/\Gamma'_i} [\Gamma_j].$$

The arrow in the bifurcation digraph is either solid, dashed, or dotted, as described in [13]. Roughly speaking, a solid arrow indicates a pitchfork bifurcation within some one-dimensional subspace of \tilde{E} , a dashed arrow indicates a transcritical bifurcation within some one-dimensional subspace of \tilde{E} , and a dotted arrow indicates a more exotic bifurcation.

In [13] we defined an *anomalous invariant subspace* (AIS) $\mathcal{A} \subseteq V$, with $V = G_M$ or H , to be a g_s -invariant subspace that is not a fixed-point subspace. Consider the PDE (1) on the cube, with f_s odd. For positive integers p , q , and r , not all 1, the set

$$(11) \quad \mathcal{A}_{p,q,r} = \text{span} \left\{ \psi_{i,j,k} \mid \frac{i}{p}, \frac{j}{q}, \frac{k}{r} \in \mathbb{Z} \right\}$$

is an AIS. The function space $\mathcal{A}_{1,1,1}$ is all of V , so it is not an AIS.

The space $\mathcal{A}_{p,q,r}$ is the appropriate function space if one solved the PDE (1) in the box $(0, \pi/p) \times (0, \pi/q) \times (0, \pi/r)$. Any solution in this box extends to a solution on the cube, with nodal planes dividing the cube into pqr boxes. These anomalous invariant subspaces are caused by the so-called hidden symmetry [10, 16] of the problem that is related to the symmetry of the PDE on all of \mathbb{R}^3 .

There are in fact a multitude of AIS for the PDE on the cube. There are proper subspaces of every AIS $\mathcal{A}_{p,q,r}$ consisting of functions with symmetry on the domain $(0, \pi/p) \times (0, \pi/q) \times (0, \pi/r)$. For example, consider the case where $p = q = r > 1$. Since there are 323 fixed-point subspaces for functions on the cube, including the zero-dimensional $\{0\}$, there are 322 AIS that are subspaces of $\mathcal{A}_{p,p,p}$, for each $p > 1$. Our code is capable of finding solutions in any AIS, within reason. The theory of anomalous solutions within AIS is unknown. The book [7] is a good reference on invariant spaces of nonlinear operators in general.

2.3. Isotypic Decomposition. To analyze the bifurcations of a branch of solutions with symmetry Γ_i , we need to understand the isotypic decomposition of the action of Γ_i on \mathbb{R}^n .

Suppose a finite group Γ acts on $V = \mathbb{R}^n$ according to the representation $g \mapsto \alpha_g : \Gamma \rightarrow \text{Aut}(V) \cong \text{GL}_n(\mathbb{R})$. In our applications we choose $\Gamma \in \mathcal{G}$ and the group action is the one in Equation (10). Let $\{\alpha_\Gamma^{(k)} : \Gamma \rightarrow \text{GL}_{d_\Gamma^{(k)}}(\mathbb{R}) \mid k \in K_\Gamma\}$ be the set of irreducible representations of Γ over \mathbb{R} , where $d_\Gamma^{(k)}$ is the dimension of the representation. We write $\alpha^{(k)}$ and K when the subscript Γ is understood. It is a standard result of representation theory that there is an orthonormal basis $B_\Gamma = \bigcup_{k \in K} B_\Gamma^{(k)}$ for V such that $B_\Gamma^{(k)} = \bigcup_{l=1}^{L_k} B_\Gamma^{(k,l)}$ and $[\alpha_g|_{V_\Gamma^{(k,l)}}]_{B_\Gamma^{(k,l)}} = \alpha^{(k)}(g)$ for all $g \in \Gamma$, where $V_\Gamma^{(k,l)} := \text{span}(B_\Gamma^{(k,l)})$. Each $V_\Gamma^{(k,l)}$ is an irreducible subspace of V . Note that $B_\Gamma^{(k)}$ might be empty for some k , corresponding to $V_\Gamma^{(k)} = \{0\}$. The *isotypic decomposition* of V under the action

of Γ is

$$(12) \quad V = \bigoplus_{k \in K} V_\Gamma^{(k)},$$

where $V_\Gamma^{(k)} = \bigoplus_{l=1}^{L_k} V_\Gamma^{(k,l)}$ are the *isotypic components*.

The isotypic decomposition of V under the action of each Γ_i is required by our algorithm. The decomposition under the action of $\text{Aut}(G)$ is the same as the decomposition under the action of Γ_0 . While there are twice as many irreducible representations of $\Gamma_0 = \text{Aut}(G) \times \mathbb{Z}_2$ as there are of $\text{Aut}(G)$, if $\alpha_{\Gamma_0}^{(k)}(-1) = I$ then $V_{\Gamma_0}^{(k)} = \{0\}$. The other half of the irreducible representations have $\alpha_{\Gamma_0}^{(k)}(-1) = -I$. The irreducible representations of Γ_0 and of $\text{Aut}(G)$ can be labeled so that $V_{\Gamma_0}^{(k)} = V_{\text{Aut}(G)}^{(k)}$ for $k \in K_{\text{Aut}(G)}$.

The isotypic components are uniquely determined, but the decomposition into irreducible spaces is not. Our goal is to find $B_\Gamma^{(k)}$ for all k by finding the projection $P_\Gamma^{(k)} : V \rightarrow V_\Gamma^{(k)}$. To do this, we first need to introduce representations over the complex numbers \mathbb{C} for two reasons. First, irreducible representations over \mathbb{C} are better understood than those over \mathbb{R} . Second, our GAP program uses the field \mathbb{C} since irreducible representations over \mathbb{R} are not readily obtainable by GAP.

There is a natural action of Γ on $W := \mathbb{C}^n$ given by the representation $g \mapsto \beta_g : \Gamma \rightarrow \text{Aut}(W)$ such that β_g and α_g have the same matrix representation. The isotypic decomposition $W = \bigoplus_{k \in \tilde{K}} W_\Gamma^{(k)}$ is defined as above using the set $\{\beta^{(k)} : \Gamma \rightarrow \text{GL}_{d_\Gamma^{(k)}}(\mathbb{C}) \mid k \in \tilde{K}_\Gamma\}$ of irreducible representations of Γ over \mathbb{C} .

The *characters* of the irreducible representation $\beta^{(k)}$ are $\chi^{(k)}(g) := \text{Tr } \beta^{(k)}(g)$. The projection $Q_\Gamma^{(k)} : W \rightarrow W_\Gamma^{(k)}$ is known to be

$$(13) \quad Q_\Gamma^{(k)} = \frac{\tilde{d}_\Gamma^{(k)}}{|\Gamma|} \sum_{g \in \Gamma} \chi^{(k)}(g) \beta_g.$$

We are going to get the $P_\Gamma^{(k)}$'s in terms of $Q_\Gamma^{(k)}$'s. A general theory for constructing these projections can be found in [13], but here it is enough that if $\chi^{(k)} = \overline{\chi^{(k)}}$, then $P_\Gamma^{(k)} = Q_\Gamma^{(k)}|_V$ and $d_\Gamma^{(k)} = \tilde{d}_\Gamma^{(k)}$, whereas if $\chi^{(k)} \neq \overline{\chi^{(k)}}$, then $P_\Gamma^{(k)} = \left(Q_\Gamma^{(k)} + \overline{Q_\Gamma^{(k)}} \right)|_V$ and $d_\Gamma^{(k)} = 2\tilde{d}_\Gamma^{(k)}$, for all $k \in K_\Gamma$.

2.4. Basis processing. In this subsection, we describe how the package of programs from [13] is modified to generate the basis needed to approximate solutions to PDE (1) on the square or cube. In principle a brute force method is possible, wherein the very large graph Ω_N is used. However, the GAP and *Mathematica* portions of the package in [13] cannot process such large graphs with current computer systems. To remedy this, we wrote specialized *Mathematica* basis-generation programs for the square and cube. These programs use the known eigenfunctions, $\psi_{i,j}$ or $\psi_{i,j,k}$, together with the GAP output from the graphs in Figure 1 to generate the data needed by the GNGA program.

It can be shown that for polynomial f_s and the known eigenfunctions ψ_m defined in terms of sine functions on the square and cube, the midpoint numerical integration can be made exact, up to the arithmetic precision used in the computation. In particular, consider the case when f_s is cubic and the eigenfunctions are $\psi_{i,j}$ (for $d = 2$) or $\psi_{i,j,k}$ (for $d = 3$). Let \tilde{M} be the desired number of sine frequencies in each dimension, that is, $i, j, k \leq \tilde{M}$. Then our midpoint numerical integration is exact if $2\tilde{M} + 1$ gridpoints in each direction are used, giving a total of $N = (2\tilde{M} + 1)^d$ total gridpoints. For example, for analyzing PDE (1) on $\Omega = (0, \pi)^2$ we used $\tilde{M} = 30$ sine frequencies and, for exact integration, $N = (2 \cdot 30 + 1)^2 = 3721$ gridpoints. Similarly, for the cube $\Omega = (0, \pi)^3$, we used $\tilde{M} = 15$ sine frequencies and $N = (2 \cdot 15 + 1)^3 = 29,791$ gridpoints.

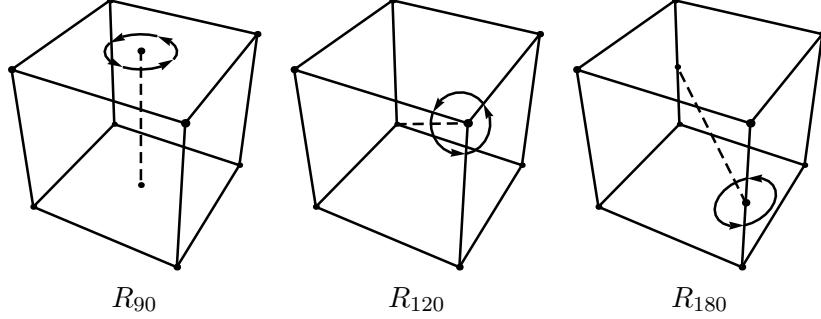


FIGURE 2. R_{90} , R_{120} , and R_{180} are matrices that rotate the cube centered at the origin by 90° , 120° and 180° , respectively. The dashed lines are the axes of rotation. The 48 element automorphism group of the cube is identified with \mathbb{O}_h . The matrix group \mathbb{O}_h is generated by these three rotation matrices and $-I_3$, the inversion through the center of the cube.

The only input for the basis generation program is \tilde{M} , the desired number of sine frequencies in each dimension. We define the bases

$$B_M = \{\psi_{i,j} \mid 1 \leq i, j \leq \tilde{M} \text{ and } \lambda_{i,j} = i^2 + j^2 < (\tilde{M} + 1)^2 + 1\}$$

for the square, and

$$B_M = \{\psi_{i,j,k} \mid 1 \leq i, j, k \leq \tilde{M} \text{ and } \lambda_{i,j,k} = i^2 + j^2 + k^2 < (\tilde{M} + 1)^2 + 2\}$$

for the cube. These bases include all eigenfunctions with eigenvalues less than $\lambda_{\tilde{M}+1,1}$ and $\lambda_{\tilde{M}+1,1,1}$ respectively. This process gives $M = 719$ for the square, not $\tilde{M}^2 = 900$, and $M = 1848$ for the cube, as opposed to $\tilde{M}^3 = 3375$.

The basis generation code also produces an automorphism file, which is an $8 \times N$ (for the square) or $48 \times N$ (for the cube) matrix describing how the elements of \mathbb{D}_4 or \mathbb{O}_h permute the vertices in Ω_N . This file is used by the GNGA program but not by the square and cube basis generation code. Instead, the action of \mathbb{D}_4 or \mathbb{O}_h on the eigenfunctions is achieved very efficiently using the standard 2D representation of \mathbb{D}_4 acting on the plane and the standard 3D representation of \mathbb{O}_h acting on \mathbb{R}^3 . This requires us to define $\psi_{i,j}$ and $\psi_{i,j,k}$ for negative integers i, j, k in this way: $\psi_{-i,j}(x, y) = \psi_{i,j}(\pi - x, y)$ and $\psi_{-i,j,k}(x, y, z) = \psi_{i,j,k}(\pi - x, y, z)$.

The GNGA program does not use the basis B_M . Rather, it uses a basis spanning the same M -dimensional space, obtained from projections of the eigenfunctions in B_M onto the isotypic components identified by the GAP program. The output of the GAP program for the small graphs shown in Figure 1 is used to achieve these projections. For the cube, the eigenspaces of the Laplacian for the eigenvalue $\lambda_{i,j,k}$ are 1-dimensional if $i = j = k$, 3-dimensional if $i = j < k$ or $i < j = k$, and 6-dimensional if $i < j < k$. Each of these eigenspaces are separated into their isotypic components automatically by the basis generation programs, using the characters found by GAP for the graphs in Figure 1. Then the projections in Equation (13) are written as 2×2 or 3×3 matrices, where β_g is the standard 2D or 3D representation of \mathbb{D}_4 or \mathbb{O}_h , respectively. With these changes, the construction of the bases proceeds as described in [13].

2.5. The symmetry of a cube. In this section we describe the various matrix groups related to the symmetry of a cube. We start with the symmetry group of the cube $(-1, 1)^3$. The matrix form of this symmetry group is the 48-element $\mathbb{O}_h := \langle R_{90}, R_{120}, R_{180}, -I_3 \rangle \leq \mathrm{GL}_3(\mathbb{R})$ where

$$R_{90} = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad R_{120} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad R_{180} = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

as shown in Figure 2, and I_3 is the 3×3 identity matrix.

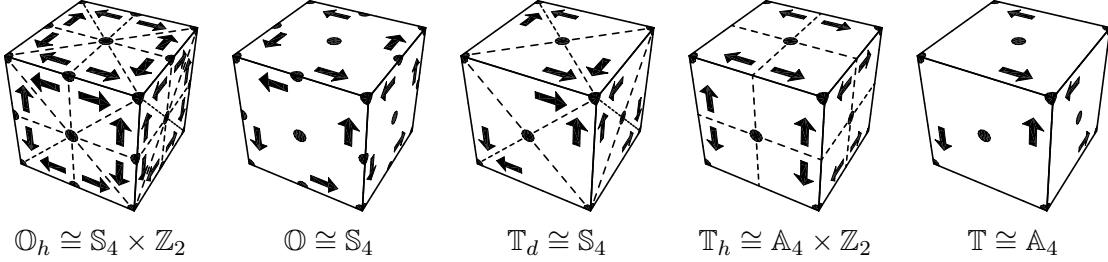


FIGURE 3. Visualization of \mathbb{O}_h and some of its subgroups. The group orbit of a single arrow is drawn, suggesting a symmetric vector field. These figures show the symmetry of the Lyapunov-Schmidt reduced gradient \tilde{g} on 3-dimensional critical eigenspaces in bifurcations with symmetry occurring in PDE (1) on the cube. The dotted lines show the intersections of reflection planes with the cube. The dots indicate 1-dimensional fixed-point subspaces of the action, which intersect the cube at a vertex, edge, or face. The EBL states that solution branches bifurcate in the direction of these 1-dimensional fixed-point subspaces in \tilde{E} .

There is a slight complication in describing the symmetry group of the cube $\Omega_3 := (0, \pi)^3$ that is the domain of our PDE, since Ω_3 is not centered at the origin. The action of $\text{Aut}(\Omega_3) \cong \mathbb{O}_h$ on Ω_3 is given by matrix multiplication about the center $d = \pi/2(1, 1, 1)$ of the cube. That is, for $\gamma \in \mathbb{O}_h$ and $x \in \Omega_3$, the action is defined as

$$\gamma \cdot x = d + \gamma(x - d).$$

The action of $\mathbb{O}_h \times \mathbb{Z}_2$ on the vector space of functions $u : \Omega_3 \rightarrow \mathbb{R}$ is given by

$$((\gamma, \beta) \cdot u)(x) = \beta u(\gamma^{-1} \cdot x), \text{ for } \gamma \in \mathbb{O}_h, \beta \in \mathbb{Z}_2 = \{1, -1\}.$$

It follows that the action of the generators of $\mathbb{O}_h \times \mathbb{Z}_2$ on eigenfunctions is

$$\begin{aligned} (R_{90}, 1) \cdot \psi_{i,j,k} &= (-1)^{j-1} \psi_{j,i,k} \\ (R_{120}, 1) \cdot \psi_{i,j,k} &= \psi_{k,i,j} \\ (R_{180}, 1) \cdot \psi_{i,j,k} &= (-1)^{k-1} \psi_{j,i,k} \\ (-I_3, 1) \cdot \psi_{i,j,k} &= (-1)^{i+j+k-1} \psi_{i,j,k} \\ (I_3, -1) \cdot \psi_{i,j,k} &= -\psi_{i,j,k}. \end{aligned}$$

There are three subgroups of \mathbb{O}_h with 24 elements. They are

$$\mathbb{O} = \langle R_{90}, R_{120}, R_{180} \rangle, \quad \mathbb{T}_d = \langle -R_{90}, R_{120}, -R_{180} \rangle, \quad \text{and } \mathbb{T}_h = \langle R_{90}^2, R_{120}, R_{180}, -I_3 \rangle$$

shown in Figure 3. The group \mathbb{O} contains the rotational symmetries of the cube (or octahedron), and is called the *octahedral group*. Note that \mathbb{O}_h is the internal direct product $\mathbb{O} \times \langle -I_3 \rangle$.

The groups \mathbb{T}_d and \mathbb{T}_h are related to the 12-element *tetrahedral group*

$$\mathbb{T} = \langle R_{90}^2, R_{120}, R_{180} \rangle,$$

that contains all of the rotational symmetries of the tetrahedron. Note that \mathbb{T}_h is the internal direct product $\mathbb{T} \times \langle -I_3 \rangle$, whereas \mathbb{T}_d does *not* contain $-I_3$. The groups \mathbb{O} and \mathbb{T}_d are both isomorphic to the symmetric group \mathbb{S}_4 . The isomorphism can be proved by considering how the groups permute the 4 diagonals of the cube. Similarly, \mathbb{T} is isomorphic to \mathbb{A}_4 , the alternating group. The group names involving \mathbb{S}_4 and \mathbb{A}_4 are used by the GAP program. In our GAP programs \mathbb{O} and \mathbb{T}_d are computed as irreducible representations of \mathbb{S}_4 . Our website, which is written automatically using the GAP output, also uses these names.

3. ALGORITHMS

The main mathematical algorithms used in the code for the results found in this paper are the tGNGA, cGNGA, and the secant method with recursive bisection, all of which are developed and described in detail in [13]. We give a brief overview of these algorithms in Section 3.1. The current implementation of C++ code which supervises the execution of these algorithms has two substantial modifications not found in [13].

The first major improvement concerns the efficient way we now use our symmetry information to reduce the matrix dimension when setting up the system for the search direction χ used in the tGNGA. In brief, each system uses only the rows and columns corresponding to the eigenfunctions having the symmetry of points on the branch. The time savings can be substantial when seeking solutions with a lot of symmetry, since the numerical integrations required to form the systems are generally the most time-intensive computations we make. We present the details in Subsection 3.2.

Secondly, the current implementation is in parallel. A serial implementation could not reproduce our results in a reasonable time. To that end, in [15] we developed a simple and easy to apply methodology for using high-level, self-submitting parallel job queues in an MPI (Message Passing Interface [6]) environment. In that paper, we apply our parallel job queue techniques toward solving computational combinatorics problems, as well as provide the necessary details for implementing our PDE algorithms in parallel C++ code in order to obtain the results found in the current article. We include a high-level description in Section 3.3.

3.1. GNGA. To follow branches and find bifurcations we take the parameter s to be the $(M + 1)^{\text{st}}$ unknown. When we say that $p = (a, s) \in \mathbb{R}^{M+1}$ is a solution we mean that $\mathbf{u} = \sum_{m=1}^M a_m \psi_m$ solves Equation (2) with parameter s , that is $g = 0$. The $(M + 1)^{\text{st}}$ equation, $\kappa(a, s) = 0$, is chosen in two different ways, depending on whether we are implementing the tangent-augmented Newton's method (tGNGA) to force the following of a tangent of a bifurcation curve or we are implementing the cylinder-augmented Newton's method (cGNGA) to force the switching to a new branch at a bifurcation point. In either case, the iteration we use is:

- compute the constraint κ , gradient vector $g := g_s(\mathbf{u})$, and Hessian matrix $h := h_s(\mathbf{u})$
- solve $\begin{bmatrix} h & \frac{\partial g}{\partial s} \\ (\nabla_a \kappa)^T & \frac{\partial \kappa}{\partial s} \end{bmatrix} \begin{bmatrix} \chi_a \\ \chi_s \end{bmatrix} = \begin{bmatrix} g \\ \kappa \end{bmatrix}$
- $(a, s) \leftarrow (a, s) - \chi$, $u = \sum a_j \psi_j$.

Equations (6) and (8) are used to compute g and h . The $(M + 1)^{\text{st}}$ row of the matrix is defined by $(\nabla_a \kappa, \frac{\partial \kappa}{\partial s}) = \nabla \kappa \in \mathbb{R}^{M+1}$; the search direction is $\chi = (\chi_a, \chi_s) \in \mathbb{R}^{M+1}$. Since this is Newton's method on $(g, \kappa) \in \mathbb{R}^{M+1}$ instead of just $g \in \mathbb{R}^M$, when the process converges we have not only that $g = 0$ (hence $p = (a, s)$ is an approximate solution to Equation (1)), but also that $\kappa = 0$.

We use the tGNGA to follow branches. The details are given in Algorithm 1 of [13]. In brief, given consecutive *old* and *current* solutions p_{old} and p_{cur} along a symmetry invariant branch, we compute the (approximate) tangent vector $v = (p_{\text{cur}} - p_{\text{old}})/\|p_{\text{cur}} - p_{\text{old}}\| \in \mathbb{R}^{M+1}$. The initial *guess* is then $p_{\text{gs}} = p_{\text{cur}} + cv$. The speed c has a minimum and maximum range, for example from 0.01 to 0.4, and is modified dynamically according to various heuristics. For example, this speed is decreased when the previous tGNGA call failed or the curvature of the branch is large, and is increased toward an allowed maximum otherwise. For the tGNGA, the constraint is that each iterate $p = (a, s)$ must lie on the hyperplane passing through the initial guess p_{gs} , perpendicular to v . That is, $\kappa(a, s) := (p - p_{\text{gs}}) \cdot v$. Easily, one sees that $(\nabla_a \kappa(a, s), \frac{\partial \kappa}{\partial s}(a, s)) = v$. In general, if f_s has the form $f_s(u) = su + H(u)$, then $\frac{\partial g}{\partial s} = -a$. Our function `tGNGA`(p_{gs}, v) returns, if successful, a new solution p_{new} satisfying the constraint. Figure 4 shows how repeated tGNGA calls are made when a worker executes a branch following job.

The constraint used by the cGNGA (see Algorithm 3 of [13]) at a bifurcation point p^* instead forces the new solution p_{new} to have a non-zero projection onto a subspace E of the critical eigenspace \tilde{E} . To ensure that we find the mother solution rather than a daughter, we insist that the Newton

```

repeat
  wait for a message from the boss
  switch the message is a
    case follow branch job
      while branch is in window do
        compute next point on branch with tGNGA
        if change in MI then
          call secant-bisection to find intervening bifurcation points
          for each intervening bifurcation point do
            L put find daughters job on queue
      use interpolated guess and tGNGA to get last point on window boundary
    case find daughters job
      for each possible bifurcation subspace of critical eigenspace do
        while new solutions still being found do
          make a random guess in subspace and call cGNGA
          if new solution is nonconjugate to previously found solutions then
            L put follow branch job on queue
    case stop command
      L stop

```

FIGURE 4. Pseudo code for the main loop of the workers. This loop is entered after loading basis and symmetry files. Whenever idle, each worker accepts and runs jobs whenever such jobs exist. The boss puts the trivial solution branch on the job queue as the first job. It manages the queue while the workers do their jobs, until the queue is empty, and then sends all workers a stop job.

iterates belong to the cylinder $C := \{(a, s) \in \mathbb{R}^{M+1} : \|P_E(a - a^*)\| = \varepsilon\}$, where P_E is the orthogonal projection onto E and the radius ε is a small fixed parameter. At a symmetry breaking bifurcation the critical eigenspace is orthogonal to the fixed-point subspace of the mother, so the mother branch does not intersect the cylinder. The constraint we use to put each Newton iterate on the cylinder is $\kappa(a, s) = \frac{1}{2}(\|P_E(a - a^*)\|^2 - \varepsilon^2) = 0$. The initial guess we use is $p_{\text{gs}} := (a^*, s^*) + \varepsilon(e, 0)$, where e is a randomly chosen unit vector in E . Clearly, p_{gs} lies on the cylinder C . A computation shows that $\nabla_a \kappa(a, s) = P_E(a - a^*)$, and $\frac{\partial \kappa}{\partial s}(a, s) = 0$. When successful, $\text{cGNGA}(p^*, p_{\text{gs}}, E)$ returns a new solution p_{new} that lies on the cylinder C .

In the above paragraph, we take E to be various low-dimensional subspaces of the critical eigenspace, corresponding to the symmetries of solutions that are predicted by bifurcation theory. For example, at an EBL bifurcation E is spanned by a single eigenvector. When the dimension of E is greater than one, we call cGNGA repeatedly with several random choices of the critical eigenvector e . The details are given in Equation (7) and Algorithm 3 of [13]. The theory we apply does not guarantee a complete prediction of all daughter solutions. Therefore we also call cGNGA with E equal to the full critical eigenspace. In this way, if the dimension of the critical eigenspace is not too big we have a high degree of confidence that we are capturing all relevant solutions, including those that arise due to accidental degeneracy and that are neither predicted nor ruled out by understood bifurcation theory. The number of guesses in each subspace is heuristically dependent on the dimension of the subspace. Too many guesses wastes time, and too few will cause bifurcating branches to be missed. Figure 4 shows how repeated cGNGA calls are made when a worker executes a *find daughters* job.

We use the secant method to find bifurcation points. In brief, when using the tGNGA to follow a solution branch and the MI changes at consecutively found solutions, say from k at the solution p_{old} to $k+\delta$ at the solution p_{cur} , we know by the continuity of $D^2 J_s$ that there exists a third, nearby solution p^* where h is not invertible and the r^{th} eigenvalue of h is zero, where $r = k + \lceil \frac{\delta}{2} \rceil$. Let $p_0 = p_{\text{old}}$, $p_1 = p_{\text{cur}}$, with β_0 and β_1 the r^{th} eigenvalues of h at the points p_0 and p_1 , respectively.

We effectively employ the vector secant method by iterating

$$\begin{aligned}\bullet \quad p_{\text{gs}} &= p_i - \frac{(p_i - p_{i-1})\beta_i}{(\beta_i - \beta_{i-1})} \\ \bullet \quad p_{i+1} &= \text{tGNGA}(p_{\text{gs}}, v)\end{aligned}$$

until the sequence (p_i) converges. The vector $v = (p_{\text{cur}} - p_{\text{old}})/\|p_{\text{cur}} - p_{\text{old}}\|$ is held fixed throughout, while the value β_i is the newly computed r^{th} eigenvalue of h at p_i . If our function `secant`($p_{\text{old}}, p_{\text{cur}}$) is successful, it returns a solution point $p^* = (a^*, s^*)$, lying between p_{old} and p_{cur} , where h has δ zero eigenvalues within some tolerance. We take the critical eigenspace \tilde{E} to be the span of the corresponding eigenvectors. If p^* is not a turning point, then it is a bifurcation point.

In fact, it is possible that several intervening bifurcation points exist. If the secant method finds a bifurcation point that has fewer than d zero Hessian eigenvalues, there must be another bifurcation point in the interval. In Figure 4, the *secant-bisection* call refers to an implementation of Algorithm 2 from [13] entitled `find_bifpoints`, whereby such an occurrence triggers a bisection and a pair of recursive calls back to itself. Upon returning, each of the one or more found bifurcation points spawns its own *find daughters* job. In turn, each time a (non-conjugate) daughter is found, a new *follow branch* job is put on the queue.

3.2. The Block Diagonal Structure of the Hessian. The majority of the computational effort for solving PDE (1) using Newton's method comes from the computation of the entries of the Hessian matrix. The time required can be drastically reduced by taking advantage of the block diagonal structure of the Hessian that follows from the isotypic decomposition of $V = \mathbb{R}^M$, the Galerkin space.

If the initial guess u has symmetry Γ_i , then the isotypic decomposition (12) of the Γ_i action on V is $V = \bigoplus_k V_{\Gamma_i}^{(k)}$, where k labels the irreducible representations of Γ_i . We assume that $k = 0$ denotes the trivial representation, so $\text{Fix}(\Gamma_i) = V_{\Gamma_i}^{(0)}$.

For any $u \in \text{Fix}(\Gamma_i)$, the symmetry of the PDE implies that the gradient $g_s(u)$ is also in $\text{Fix}(\Gamma_i)$, and the Hessian, evaluated at u , maps each of the isotypic components to itself. That is,

$$u \in \text{Fix}(\Gamma_i) \implies g_s(u) \in \text{Fix}(\Gamma_i) \text{ and } h_s(u) \left(V_{\Gamma_i}^{(k)} \right) \subseteq V_{\Gamma_i}^{(k)}.$$

Thus, the Hessian is block diagonal in the basis B_{Γ_i} defined in Section 2.3. A huge speedup of our program is obtained by only computing the Hessian restricted to $\text{Fix}(\Gamma_i)$ when doing Newton's method. After a solution is found, the block diagonal structure of the Hessian allows its efficient computation by avoiding integration of zero terms. The full Hessian is required for the calculation of the MI.

Actually, the way we achieve a speedup in our numerical algorithm is not quite this simple. In our implementation of the GNGA we always use B_{Γ_0} , a basis of eigenvectors of the Laplacian. The basis vectors are partitioned into bases $B_{\Gamma_0}^{(k)}$ for each of the isotypic components of the Γ_0 action on V . We do not change the basis depending on the symmetry of the solution we are approximating. Hence, we do not simply compute the blocks of the block diagonal Hessian.

When doing Newton's method, we use a *reduced Hessian* \bar{h}_s in place of the full Hessian h_s . Define $P^{(k)}v$ to be the projection of v onto $V_{\Gamma_i}^{(k)}$. For $u \in \text{Fix}(\Gamma_i)$, the reduced Hessian is defined to be

$$\bar{h}_s(u)_{j,k} = \begin{cases} h_s(u)_{j,k} & \text{if } P^{(0)}\psi_j \neq 0 \text{ and } P^{(0)}\psi_k \neq 0 \\ \lambda_j - s & \text{if } j = k \text{ and } P^{(0)}\psi_j = 0 \\ 0 & \text{if } j \neq k \text{ and } (P^{(0)}\psi_j = 0 \text{ or } P^{(0)}\psi_k = 0). \end{cases}$$

For each $u \in \text{Fix}(\Gamma_i)$, assuming h_s and \bar{h}_s are nonsingular, the Newton search direction is the solution to either system $h_s(u)\chi = g_s(u)$ or $\bar{h}_s(u)\chi = g_s(u)$. The terms in the reduced Hessian of the form $(\lambda_j - s)\delta_{j,k}$ are included to make \bar{h}_s nonsingular. They are not strictly necessary since we find the least square solution χ with the smallest norm, but they improve the performance of the LAPACK solver (dgels).

After the solution is found, we identify those elements of the full Hessian that are known to be zero, and avoid doing numerical integration for those elements. In particular, $h_s(u)_{p,q} = 0$ if there is no k such that $P^{(k)}\psi_p \neq 0$ and $P^{(k)}\psi_q \neq 0$.

As a simple example for demonstration purposes, suppose that there are $M = 5$ modes, and we are using the basis vectors ψ_1, \dots, ψ_5 . Suppose further that the isotypic decomposition of the Γ_i action is $V_{\Gamma_i}^{(0)} = \text{span}\{\psi_1, \psi_3\}$, $V_{\Gamma_i}^{(1)} = \text{span}\{\psi_2\}$, $V_{\Gamma_i}^{(2)} = \text{span}\{\psi_4 + \psi_5\}$, and $V_{\Gamma_i}^{(3)} = \text{span}\{\psi_4 - \psi_5\}$. The only nonzero projections are $P^{(0)}\psi_1, P^{(0)}\psi_3, P^{(1)}\psi_2, P^{(2)}\psi_4, P^{(2)}\psi_5, P^{(3)}\psi_4$, and $P^{(3)}\psi_5$. Thus, if $u \in \text{Fix}(\Gamma_i)$, the gradient and Hessian have the form

$$(14) \quad g_s(u) = \begin{bmatrix} * \\ 0 \\ * \\ 0 \\ 0 \end{bmatrix}, \quad h_s(u) = \begin{bmatrix} * & 0 & * & 0 & 0 \\ 0 & * & 0 & 0 & 0 \\ * & 0 & * & 0 & 0 \\ 0 & 0 & 0 & * & * \\ 0 & 0 & 0 & * & * \end{bmatrix}.$$

Note that a change of basis could be done to diagonalize the lower right 2×2 block, but our program does not do this. When we solve for the search direction χ in Newton's method, we use the restricted Hessian

$$\bar{h}_s(u) = \begin{bmatrix} * & 0 & * & 0 & 0 \\ 0 & \lambda_2 - s & 0 & 0 & 0 \\ * & 0 & * & 0 & 0 \\ 0 & 0 & 0 & \lambda_4 - s & 0 \\ 0 & 0 & 0 & 0 & \lambda_5 - s \end{bmatrix},$$

where only the * terms are computed using numerical integration. Thus, in this small example, four numerical integrations are needed to compute the reduced Hessian for each step of Newton's method, and nine numerical integrations are needed to compute the full Hessian.

The speedup obtained in this manner is quite dramatic for solutions with high symmetry when M is large. As an example, for the PDE on the cube the primary branch that bifurcates at $s = 3$ has symmetry $\Gamma_2 \cong \mathbb{O}_h$. The Γ_2 action on V has 10 isotypic components. Our program, with $\tilde{M} = 15$ and hence $M = 1848$, processes these solutions 10 to 15 times faster than it processes solutions with trivial symmetry, where there is just one isotypic component and all $M(M+1)/2 = 1,708,474$ upper-triangular elements of the Hessian need to be computed for every step of Newton's method. Each one of these Hessian elements requires a sum over $N = 29,729$ grid points (see Equation (8)). Considering the whole process of finding a solution with Γ_2 symmetry using 3 or 4 steps of Newton's method, and then computing the MI of this solution, the majority of the computational effort comes from the numerical integrations needed to compute the full Hessian once after Newton's method converges. Even with the speedup indicated in (14), the single computation of $h_s(u)$ takes longer than the total time to perform all the other computations. These include the computation of the reduced Hessian $\bar{h}_s(u)$ and gradient $g_s(u)$, and the LAPACK calls to solve $\bar{h}_s(u)\chi = g_s(u)$ at each Newton step, as well as the LAPACK computation of the eigenvalues of the full Hessian evaluated at the solution.

3.3. Parallel Implementation of Branch Following. It was necessary to implement our code in parallel in order to get timely results. We use a library of functions (MPQueue) presented in [15] which make it easy to create and manage a parallel job queue. For the current application of creating a bifurcation diagram, we choose a natural way to decompose the task into two types of jobs, namely *branch following* and *find daughters* jobs. The boss starts MPI and puts the trivial solution branch on the queue as the first job. After initialization and whenever idle, each worker accepts and runs jobs whenever they exist. The boss manages the queue while the workers do their

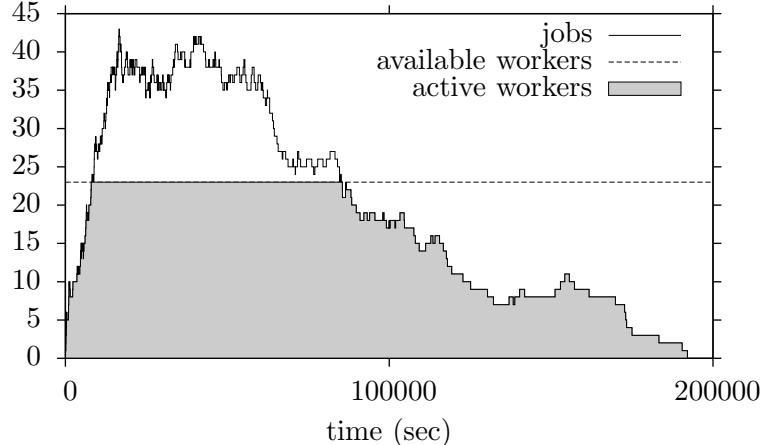


FIGURE 5. The load diagram showing the number of jobs (branch following and bifurcation processing) and the number of active workers. There were 24 processors and thus 23 workers. The jobs above the dotted line are waiting in the job queue. The total time for the run was 192,094 seconds \approx 53.4 hours. The average number of active workers is 15.6. Thus, the run would have taken approximately 34 days with a single processor.

jobs, until the queue is empty, and then sends all workers a stop command. Figure 4 shows how a parallel job queue is used to supervise the running of the jobs.

Normal termination for branch following jobs occurs when the branch exits a parameter interval for s . When a change in MI is detected during branch following, the secant method is used to find an intervening bifurcation point with a proscribed zero eigenvalue of the Hessian of J . Recursive bisection is used if the number of zero eigenvalues of a found bifurcation point does not equal the observed change in MI over a given subinterval, ensuring that all intervening bifurcation points are found. Each bifurcation point spawns a *find daughters* job. These jobs invoke the cGNGA with multiple random guesses from all possible bifurcation subspaces of the critical eigenspace. The number of attempts depends heuristically on the dimension of each subspace, stopping after sufficiently many tries have failed to find a new solution not in the group orbit of any previously found solution. Each new, bifurcating solution spawns a *follow branch* job.

We conclude this subsection with an example demonstrating the parallel run times encountered when generating our results for the cube region. In particular, Figure 5 shows the load diagram of a run for PDE (1) on the cube with $\tilde{M} = 15$, which gives $M = 1848$ modes and uses $N = 31^3$ gridpoints. We followed all branches connected to the trivial branch with $0 \leq s \leq 13$, up to 4 levels of bifurcation. The code found 3168 solutions lying on 111 non-conjugate branches, with 126 bifurcation (or fold) points. This 2-day run used 24 processors and generated all the data necessary for Figures 14, 15, 20, 21, and 22. In Table 6.1 of [15] we provide runtime data for a similar problem using a varying number of processors, in order to demonstrate scalability.

3.4. Creating contour plots. We now describe the process of creating a contour plot for visually depicting an approximate solution on the square and cube. For the square contour plot graphics found in Figures 9 and 11, we use $\tilde{M} = 30$ leading to $N = 61^2$ grid points. This is a sufficient number of grid points to generate good contour plots. The cube solutions visualized in the contour plot graphics found in Figures 15, 19, 22, 21, and 25 were all obtained via GNGA using $\tilde{M} = 15$ leading to $M = 1848$ modes and $N = 31^3$ grid points. This resolution is not quite sufficient for generating clear contour plots. Using the M coefficients and the exact known basis functions, we reconstruct u on a grid with 56^3 points before applying the techniques described below.

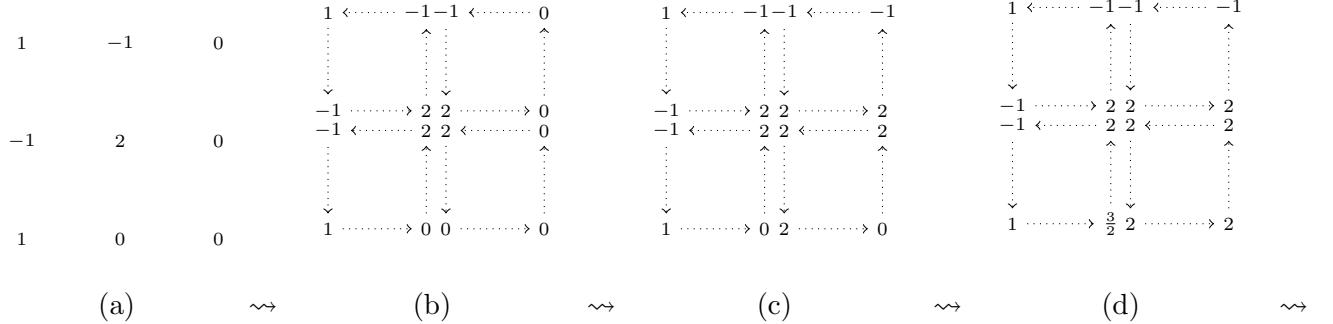


FIGURE 6. Square creation, followed by zero localization and zero purge.

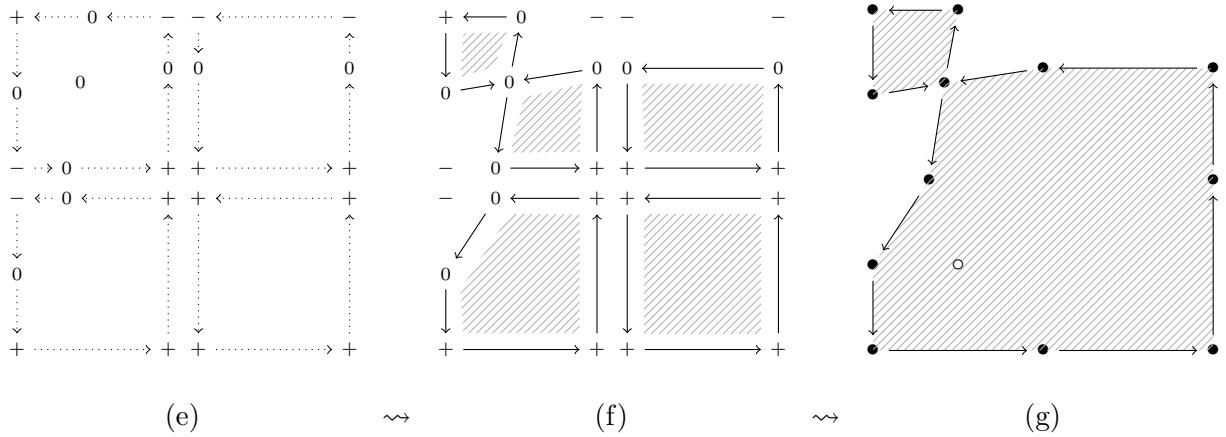


FIGURE 7. Boundary detection, followed by polygon creation and polygon merging.

When Ω is the cube, since $u = 0$ on the boundary we plot u on the boundary of a smaller concentric cube. The faces of this inner cube are 5 gridpoints deep from the original boundary. Thus each of the 6 faces we draw has u values on $(56 - 2 \cdot 5)^2$ grid points. The side length of the cube seen in the figures is about $46/56 = 82\%$ the length of the original cube. At the end of the process, the contour plots on the six faces of the inner cube are assembled into a front and back view. We use a standard perspective mapping from \mathbb{R}^3 to \mathbb{R}^2 to join three of the six faces for each view.

In the remainder of this sub-section, we describe how to make a contour map for a function $u : \Omega \rightarrow \mathbb{R}$ on the square or cube. First of all, an algorithm chooses one function in the group orbit $\Gamma_0 \cdot u$ that makes the symmetry of the contour map agree with the flag diagram. It also replaces u by $-u$ if this saves ink. For example, if the function has a 3-fold symmetry about a diagonal of the cube, our algorithm chooses a function that satisfies $(\mathbb{R}_{120}, 1) \cdot u = u$ since the axis of rotation of R_{120} points toward the viewer. Then we have the values of u on a square grid (from the square or a face of the cube) as shown in Figure 6(a). We do some data pre-processing, depicted in Figures 6(b–d), in order to get acceptable plots when there are many data points with $u = 0$, which is common because of the symmetry of solutions. After interpolation, Figure 7(e), the positive region is the union of many polygons, mostly squares, Figure 7(f). Redundant edges are eliminated to obtain a few many-sided polygons that are shaded, Figure 7(g). The contours are obtained by the same algorithm: It finds the many-sided polygons enclosing the regions where $u \geq c$, but just the boundary is drawn. Our implementation directly creates postscript files. Without the elimination of redundant edges, the files would be of an unreasonably large size, and the contours could not be drawn with the same algorithm.

We now provide the details. Let (\mathbf{x}, a) represent a point where $\mathbf{x} \in \mathbb{R}^n$, $n = 2$ or 3 , is a grid point and $a \in \mathbb{R}$ such that $u(\mathbf{x}) \approx a$ for a given solution u . We use the following steps as visualized by Figures 6 and 7.

- (1) Square creation (a)~(b): In the first step, we break the square grid into individual squares to be handled separately. We represent one of these squares with the cycle $((\mathbf{x}_1, a_1) \cdots (\mathbf{x}_4, a_4))$ containing the corner points.
- (2) Zero localization (b)~(c): We eliminate several zero values in a row by applying the rule

$$\cdots (\mathbf{x}, a)(\mathbf{y}_1, 0) \cdots (\mathbf{y}_n, 0)(\mathbf{z}, b) \cdots \rightsquigarrow \cdots (\mathbf{x}, a)(\mathbf{y}_1, a) \cdots (\mathbf{y}_n, b)(\mathbf{z}, b) \cdots$$

to each cycle where the variables a and b represent non-zero values.

- (3) Zero purge (c)~(d): We eliminate the zero values without a sign change by applying the rule

$$\cdots (\mathbf{x}, a)(\mathbf{y}, 0)(\mathbf{z}, b) \cdots \rightsquigarrow \cdots (\mathbf{x}, a)(\mathbf{y}, (a+b)/2)(\mathbf{z}, b) \cdots$$

where $ab > 0$.

- (4) Boundary detection (d)~(e): We insert zero points between sign changes using the rule

$$\cdots (\mathbf{x}, a)(\mathbf{y}, b) \cdots \rightsquigarrow \cdots (\mathbf{x}, a) \left(\frac{|b|\mathbf{x} + |a|\mathbf{y}}{|a| + |b|}, 0 \right) (\mathbf{y}, b) \cdots$$

where $ab < 0$. The location of the new zero point is determined using linear interpolation. The number k of zeros in the cycle after this step must be 0, 2 or 4.

- (5) Polygon creation (e)~(f): Now we create polygons with shaded inside to indicate positive values.

- (a) If $k = 0$ and $a_i > 0$ for all i in the cycle $((\mathbf{x}_1, a_1) \cdots (\mathbf{x}_4, a_4))$, then we create the polygon with these corner points. This occurs in the lower-right square of Figure 7(e). Note that if $a_i < 0$ then no polygon is created.
- (b) If $k = 2$, then we create a polygon for each cyclic pattern of the form

$$(\mathbf{x}, 0)(\mathbf{y}_1, a_1) \cdots (\mathbf{y}_n, a_n)(\mathbf{z}, 0)$$

with a_1, \dots, a_n all positive. There is an example of this with $n = 2$ in the upper-right square and an example with $n = 3$ in the lower-left square of Figure 7(e). It is also possible to have $n = 1$, but this is not shown in the example.

- (c) If $k = 4$, then we find the intersection $(\mathbf{c}, 0)$ of the line segments joining opposite zeros. Then we look for cyclic patterns of the form

$$(\mathbf{x}, 0)(\mathbf{y}, a)(\mathbf{z}, 0)$$

with a positive. For each such pattern we create the polygon with corner points

$$(\mathbf{x}, 0), (\mathbf{y}, a), (\mathbf{z}, 0), (\mathbf{c}, 0).$$

This happens in the upper-left square of Figure 7(e)

- (6) Polygon merging (f)~(g): We combine polygons sharing a side into a single polygon. This gives the zero set which is the boundary of the shaded positive region. This single polygon is written directly to a postscript file. Polygon merging is a time consuming operation, but writing the merged polygons instead of the individual ones reduces the size of the postscript file significantly.
- (7) Level curve creation: Level curves are created using the previous steps. We use the zero set of $u - c$ as the level curve of $u = c$. The original grid values are changed by the shift value c and only the final merged polygon is written to the postscript file without shading. Note that polygon merging is essential for finding level curves.
- (8) Local extremum dots (g): White dots are drawn at the estimated position of local extrema in the $u > 0$ region, as indicated in Figure 7(g). Every interior data point (\mathbf{x}, a) , is checked to see if a is at least as large (or at least as small) as b_i for the eight neighbors (\mathbf{x}_i, b_i) . For the data in Figure 6(a), only the center point is checked, and it is indeed a local maximum

grid point. For each local extremum grid point, a more precise estimate of the extreme point is found by fitting a quadratic function to the 9 data points with the extreme grid point at the center. Since a quadratic function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ has 6 constants, the least squares solution for the 9 equations in 6 unknowns is found. Then standard calculus gives the position of the extreme point of f , and a dot is drawn there. For the data in 6(a), the maximum point of the fitted f is located half way to the grid point to its lower left. Similarly, black dots are drawn at the estimated positions of the local extrema in the white ($u < 0$) region.

4. THE PDE ON THE SQUARE AND \mathbb{D}_4 SYMMETRY

In this section we re-visit the case where Ω is the unit square. We first studied the square case in [16], without the benefit of automation or our recent improvements in branch following. In that paper, all of the symmetry analysis was done by hand. Each of the few bifurcation points we analyzed were essentially small projects in themselves, as were monotonic segments of the branches connecting them. In the present article, we use our automatically generated bifurcation digraph and new Newton algorithms to quickly reproduce the old results and then go much further ($s \gg \lambda_1$) in finding solutions of every possible symmetry type. We also show that we can handle some accidental degeneracies that occur in the square case.

Figure 8 contains the condensed bifurcation digraph for PDE (1) on a region with \mathbb{D}_4 symmetry. This digraph, along with the files required by our GNGA code for following bifurcations, was automatically generated by our code. For clarity, we have chosen here to annotate the vertices of the digraph with schematic diagrams, i.e., contour plots of step functions on the square which display the proper symmetries in a visually obvious way. A solid line in the schematic diagrams represents a nodal line, whereas a dashed line is a line of reflectional symmetry and a dot is a center of rotational symmetry. Contour plots of actual solutions can be found in Figure 9. The bifurcation digraph was described in Subsection 2.2 and a more thorough discussion is in [13].

The bifurcation digraph in Figure 8 is condensed, as described in [13]. The symmetry types are grouped into condensation classes, and not all of the arrows are drawn. For example, one condensation class is the block of four symmetry types near the top. The condensation class has three arrows emanating from it, but in the un-condensed bifurcation digraph there are 5 arrows emanating from *each* of the 4 solution types in the block. The little numbers near the arrow tails count the number of arrows emanating from each solution. Similarly, the little numbers near the arrow heads count the number of arrows ending at each of the solution types in the condensation class.

At the top of Figure 8 is the trivial function $u \equiv 0$ whose symmetry is all of $\mathbb{D}_4 \times \mathbb{Z}_2$. At the bottom is a function with trivial symmetry, i.e., whose symmetry is the group containing only the identity. There are four generic bifurcations with \mathbb{Z}_2 symmetry from the trivial branch. For these bifurcations the critical eigenspace \tilde{E} is the one-dimensional irreducible subspace for \mathbb{Z}_2 and there is a pitchfork bifurcation creating two solution branches, (u_s, s) and $(-u_s, s)$ at some point $(0, s^*)$. The figure at the bottom left indicates the symmetry of a vector field in \tilde{E} . We can think of this as the ODE on the one-dimensional center manifold, or Lyapunov-Schmidt reduced bifurcation equation $\tilde{g} = 0$. There is one generic bifurcation with \mathbb{D}_4 symmetry from the trivial branch. Here the two-dimensional critical eigenspace \tilde{E} has lines of reflection symmetry across “edges” and non-conjugate lines of reflection symmetry across “vertices”, as shown in the bottom middle part of Figure 8. At the bifurcation there are two conjugacy classes of solution branches, therefore the bifurcation digraph has two arrows labeled \mathbb{D}_4 coming out of the trivial solution. On the condensed bifurcation digraph these two arrows are collapsed into one. Note that there are several more bifurcations with \mathbb{Z}_2 symmetry or with \mathbb{D}_4 symmetry in the bifurcation digraph. For example, each of the 4 solution types in the second row can have a bifurcation with \mathbb{D}_4 symmetry.

There is only one more type of generic bifurcation that occurs in PDE (1) on the square: a bifurcation with \mathbb{Z}_4 symmetry. This is a generic bifurcation in this *gradient* system [13], and

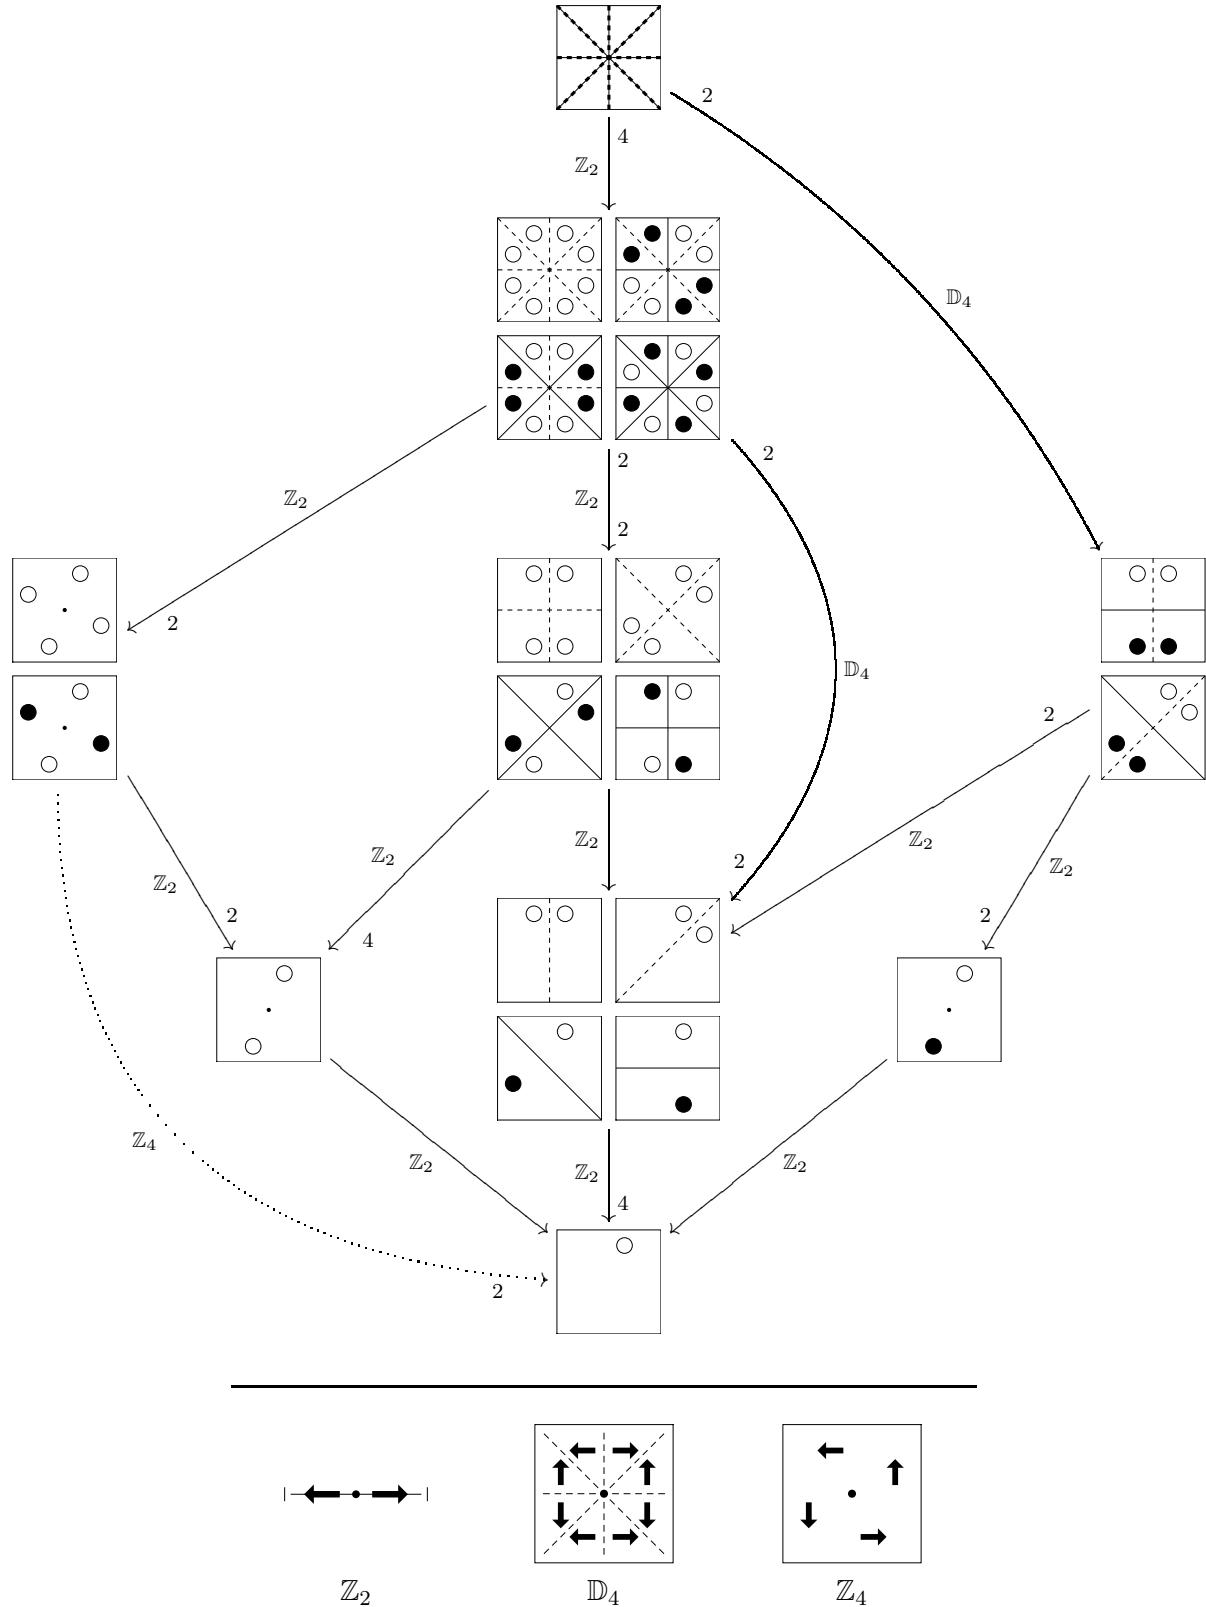


FIGURE 8. Condensed bifurcation digraph for PDE on the square, and the irreducible spaces for the generic bifurcations in the digraph.

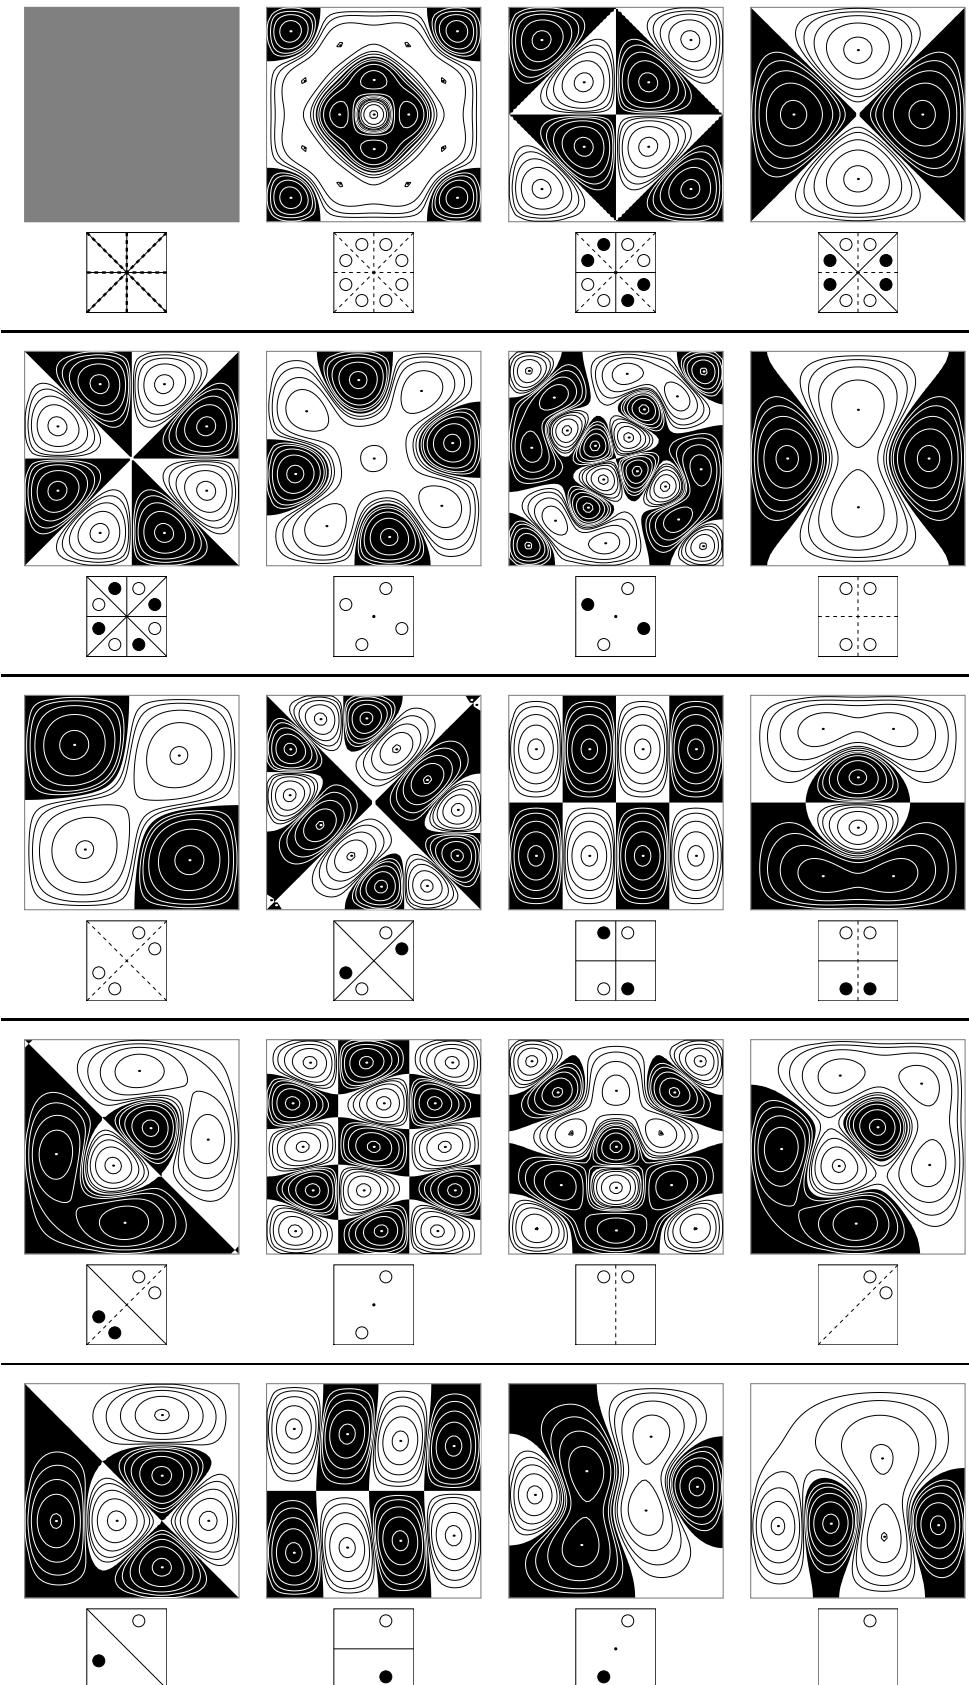


FIGURE 9. Contour plots of solutions on the square of each symmetry type, with schematic diagrams (see Figure 8).

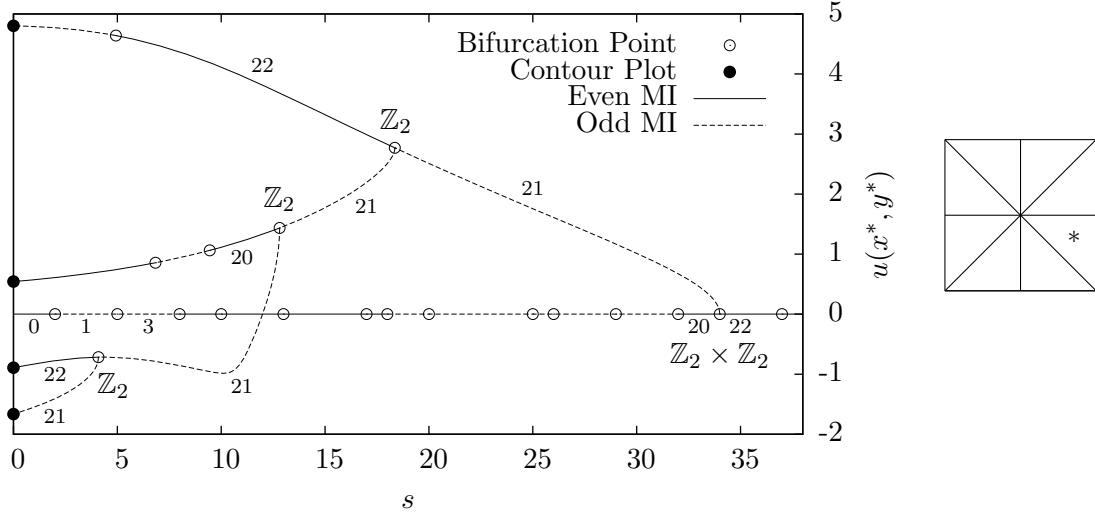


FIGURE 10. A partial bifurcation diagram for PDE (1) on the square, showing just one primary branch and selected daughter branches. The value of $u(x^*, y^*)$ vs. s is plotted, where (x^*, y^*) is a *generic point* of the square, as shown in the figure on the right. A generic point is not on any of the lines of reflection symmetry, and we choose a point equidistant from those lines and the boundary. The open dots show bifurcation points, and the group of the bifurcation is shown for some. The small numbers indicate the Morse Index of the solutions, and solid lines show an even MI whereas dashed lines show an odd MI. The solid dots at $s = 0$, from top to bottom, correspond to the contour plots in Figure 11, from left to right.

the daughter solutions can be anywhere in the two-dimensional irreducible subspace (except the origin) since there are no lines of reflection symmetry. Note the use of a dotted arrow type for this bifurcation, and that there are no dashed arrows in this particular bifurcation digraph. The reduced bifurcation equations on the critical eigenspace at this bifurcation has the symmetry indicated at the lower right part of Figure 8.

Figure 9 contains the contour plot of an example solution to Equation (1) at $s = 0$ for each of the 20 possible symmetry types on the square. The contour heights are $\pm c2^{-h}$ with $h \in \{0, \dots, 4\}$ and an appropriate c near $\max(|u|)$, to give more contours near $u = 0$. These figures were made with $\tilde{M} = 30$, meaning that the largest frequency in each direction is 30. Thus the mode with the smallest eigenvalue that is left out of the basis is $\psi_{31,1}$. This leaves $M = 719$ modes in our basis. Unlike [16], where an initial guess for each branch needed to be input by humans, the solutions in Figure 9 were found automatically by following all of the primary branches that bifurcated from the trivial solution, and all of the secondary branches, etc., recursively as described in [13].

In addition to generic bifurcations, the PDE on the square has degenerate bifurcations due to the “hidden symmetry” of translation in the space of periodic functions [10]. For example, the bifurcation point at $s = \lambda_{3,5} = 34$ on the trivial branch $a^* = 0$ has an accidental degeneracy of Type 2 as defined in [13]. Figure 10 shows a partial bifurcation diagram containing this point and 3 levels of branches bifurcating from it; corresponding contour plots of solutions are found in Figure 11. A bifurcation diagram showing branches that bifurcate at $s \leq \lambda_{2,3} = 13$ is shown in [16].

The 2-dimensional critical eigenspace at this primary bifurcation point is $\tilde{E} = \text{span}\{\psi_{3,5}, \psi_{5,3}\}$. The trivial branch, whose symmetry is $\Gamma_0 \cong \mathbb{D}_4 \times \mathbb{Z}_2$, undergoes a bifurcation with $\Gamma_0/\Gamma'_0 \cong \mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry at $s = 34$. The action of $\mathbb{Z}_2 \times \mathbb{Z}_2$ on \tilde{E} is generated by

$$b\psi_{3,5} + c\psi_{5,3} \mapsto c\psi_{3,5} + b\psi_{5,3} \quad \text{and} \quad b\psi_{3,5} + c\psi_{5,3} \mapsto -b\psi_{3,5} - c\psi_{5,3}.$$

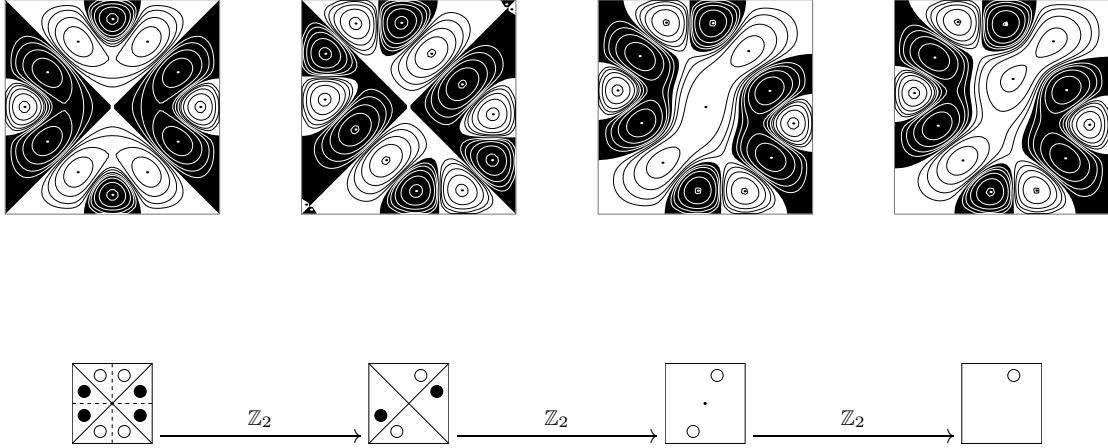


FIGURE 11. A sequence of solutions obtained at $s = 0$ by following a chain of \mathbb{Z}_2 bifurcations from the trivial solution with full $\mathbb{D}_4 \times \mathbb{Z}_2$ symmetry, down to a solution with trivial symmetry. The sequence of bifurcations shown is a path in the bifurcation digraph in Figure 8. The first solution shown is on a primary branch bifurcating at $s = \lambda_{3,5} = \lambda_{5,3} = 34$, and is approximately a multiple of $\psi_{3,5} - \psi_{5,3}$. The four contour plots represent the solutions indicated by the four black dots in Figure 10.

Our code uses the ordered basis $(\psi_{3,5} + \psi_{5,3}, \psi_{3,5} - \psi_{5,3})$. In this basis, the action of Γ_0/Γ'_0 on \tilde{E} is isomorphic to the natural action of

$$\langle \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} \rangle \cong \mathbb{Z}_2 \times \mathbb{Z}_2$$

on $[\tilde{E}] = \mathbb{R}^2 = \mathbb{R} \oplus \mathbb{R}$. Note that \tilde{E} is not an irreducible space; this is a degenerate bifurcation. The basis vectors were chosen to span the two one-dimensional irreducible subspaces, which are also fixed-point subspaces of the Γ_0/Γ'_0 action on \tilde{E} and therefore \tilde{g} -invariant subspaces. Each of these subspaces is used as an $E \subseteq \tilde{E}$ in the cGNGA algorithm described in Section 3.1. A pitchfork bifurcation generically occurs in each of these one-dimensional invariant spaces.

The one-dimensional subspaces $\text{span}\{\psi_{3,5}\}$ and $\text{span}\{\psi_{5,3}\} \subseteq \tilde{E}$, corresponding to $\text{span}\{(1, 1)\}$ and $\text{span}\{(-1, 1)\} \subseteq \mathbb{R}^2$, respectively, are not fixed-point subspaces. However, they are AIS (see Section 2.2) of $\tilde{g} : \tilde{E} \rightarrow \tilde{E}$, since $\psi_{3,5}$ (and $\psi_{5,3}$) can be periodically extended to tile the plane with a solution to the PDE (1). There is a primary branch of solutions which is tangent to $\{(u, s) = (a\psi_{3,5}, 34) \mid a \in \mathbb{R}\}$ at $(0, 34)$. Thus, there are at least three (conjugacy classes of) solution branches bifurcating from this degenerate bifurcation with $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry. Near the bifurcation, the nontrivial solutions are approximately multiples of $\psi_{3,5} + \psi_{5,3}$, $\psi_{3,5} - \psi_{5,3}$, or $\psi_{3,5}$ (or its conjugate $\psi_{5,3}$). Figure 10 shows a partial bifurcation diagram which follows the primary branch which is approximately a multiple of $\psi_{3,5} - \psi_{5,3}$ near the bifurcation.

Figure 11 shows contour plots of example solutions along a particular path in the bifurcation digraph shown in Figure 8. The primary branch is created at the degenerate bifurcation with $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry at $s = 34$ discussed above. The critical eigenspace is two-dimensional at $s = \lambda_{3,5} = 34$, and this bifurcation is not on the bifurcation digraph, Figure 8, which only shows generic bifurcations. There are two primary branches conjugate to the one shown, four secondary branches, eight tertiary, and 16 branches conjugate to the solution with trivial symmetry shown in Figure 11. At each bifurcation our GNGA code follows exactly one of the conjugate branches that bifurcate.

Figure 12 contains a numerical demonstration of the convergence of the GNGA as the number of modes increases. In Figure 7(b) of [16], we previously provided a portion of a similar graph of the L^2 norm of $\Delta u + u^3$ vs. \tilde{M} , for a particular solution u . That graphic was not entirely

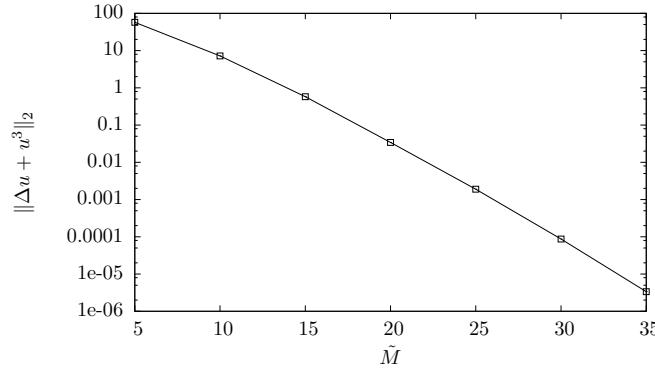


FIGURE 12. A measure of the error of our approximation to the CCN solution (minimal energy sign-changing solution [4]) of PDE (1) at $s = 0$, as a function of \tilde{M} . The region used was $\Omega = (0, 1)^2$ to facilitate comparison with Figure 7(b) of [16].

convincing in showing convergence to 0. Using our current code with larger values of \tilde{M} and a smaller convergence tolerance in Newton’s Method, we re-computed the data for the same problem and obtained the more accurate numerical result displayed here.

5. THE PDE ON THE CUBE

The PDE (1) on the cube has a rich array of bifurcations with symmetry. If Ω is a planar region, only bifurcations with \mathbb{D}_n or \mathbb{Z}_n symmetry are present in the bifurcation digraph. The bifurcation digraph of the PDE on the cube includes bifurcations with \mathbb{Z}_2 symmetry, which have 1-dimensional critical eigenspaces. The digraph also includes bifurcations with \mathbb{Z}_n or \mathbb{D}_n symmetry, $n \in \{3, 4, 6\}$, for which the critical eigenspace is 2-dimensional. A novel feature of the PDE on the cube is that there are 5 bifurcations with symmetry that have 3-dimensional critical eigenspaces, shown in Figure 3.

Recall [13] that a generic *bifurcation with Γ symmetry* has a critical eigenspace \tilde{E} that is a faithful, irreducible representation space of Γ . Faithful means that only the identity in Γ acts trivially on \tilde{E} , and irreducible means that no proper subspace of \tilde{E} is Γ -invariant.

Faithful, irreducible representation spaces for \mathbb{D}_4 and \mathbb{Z}_4 were shown in Figure 8, and the generalization to \mathbb{D}_n and \mathbb{Z}_n is obvious. Note that there are no such representation spaces for $\mathbb{Z}_2 \times \mathbb{Z}_2$, since no 1-dimensional representation space is faithful, and every 2 or higher-dimensional representation space is reducible. Recall that the bifurcation with $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry that occurs at $s = 34$ in Figure 10 is not *generic*. That bifurcation point has a Type-2 degeneracy [13].

This section contains our main new numerical results, namely approximate solutions to Equation (1) on the cube. For convenience, we denote the 99 symmetry types of solutions to this PDE on this region by S_0, \dots, S_{98} . Obtaining accurate approximations on this 3-dimensional region requires a large number of gridpoints. We use a parallel implementation as described in Section 3.3.

In Subsection 5.1, we give an overview of features of the bifurcation digraph, which is too big to include in its entirety in a single document. We describe how our companion website [14] can be used to navigate the digraph in order to view graphics and understand various symmetry information across the spectrum of solutions. In the remaining subsections, we include a survey of our numerical results which showcase our analysis of the bifurcations with most interesting symmetries.

5.1. The Bifurcation Digraph. The bifurcation digraph of the $\mathbb{O}_h \times \mathbb{Z}_2$ action on $V = G_M$ or H is far too complicated to display as a figure in this paper. Our web site [14] has a page for each symmetry type S_i , encoding the arrows emanating from this symmetry type, along with additional information.

Symmetry types, partitioned into condensation classes.
Non-isomorphic types are separated with a ‘:’
:0: 1 2 3 4 : ⋯ : 19 : 20 21 22 23 : ⋯ : 87 : 88 . 89 90 91 92 . 93 94 . 95 . 96 97 : 98
Bifurcates from: 0 2
Symmetry type: 21, representatives are isomorphic to \mathbb{D}_6 (Flag diagram and back and front contour plots appear here.)
Bifurcation with (symmetry) to (symmetry type): $\mathbb{Z}_2 \rightarrow 55$ $\mathbb{Z}_2 \rightarrow 58$ $\mathbb{Z}_2 \rightarrow 57$ $\mathbb{D}_6 \rightarrow 91 92$ $\mathbb{D}_3 - \rightarrow 80$
View: 55 58 57 91 92 80

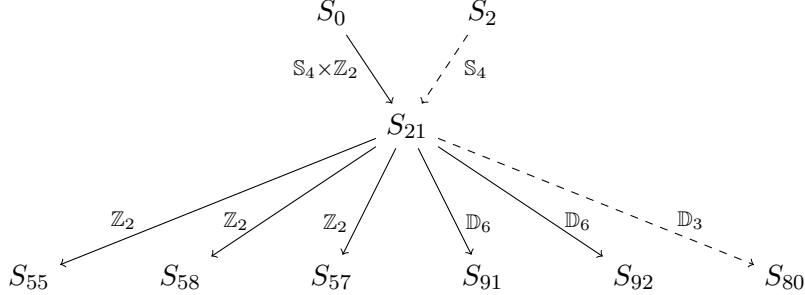


FIGURE 13. A schematic representation of the page for symmetry type S_{21} from the companion web site [14], along with the corresponding arrows in the bifurcation digraph. The labels on the top two arrows are found in the pages for symmetry types S_0 and S_2 , respectively.

An example from the companion web site for symmetry type S_{21} is shown in the top half of Figure 13. The first box contains links to all the symmetry type pages. The 99 symmetry types are grouped into isomorphism classes by colons. The isomorphism classes are further subdivided into condensation classes by periods. The abbreviated list in Figure 13 shows, for example, that S_0 , S_{19} , S_{87} and S_{99} are singleton condensation classes, and that $\{S_{20}, S_{21}, S_{22}, S_{23}\}$ is a condensation class. The symmetry types S_{88} through S_{97} are all isomorphic. These 11 symmetry types are separated into 5 condensation classes. For example $\{S_{93}, S_{94}\}$ is a condensation class.

The second box indicates that there are arrows in the bifurcation digraph pointing from S_0 and from S_2 to S_{21} .

The third box indicates that any $\Gamma_i \in S_{21}$ is isomorphic to \mathbb{D}_6 , and contains the graphics on the web page. There is a flag diagram for every symmetry type, and contour plots if our computer program found a solution with this symmetry type. About half of the symmetry types feature contour plots.

The fourth box encodes the 6 arrows in the bifurcation digraph emanating from S_{21} , as shown in the bottom half of Figure 13. In addition, the arrows are separated into the 5 generic bifurcations with symmetry coming from the 5 nontrivial irreducible representations of \mathbb{D}_6 . The symmetry types in any condensation class have an identical pattern of generic bifurcations, with different labels of the symmetry types. Thus, there are 6 arrows emanating from each of the symmetry types in the class $\{S_{20}, S_{21}, S_{22}, S_{23}\}$. In the condensed bifurcation digraph, these 24 arrows are represented by just 6 arrows.

The final box contains buttons to view selected daughter flag diagrams.

5.2. Bifurcations from the first three eigenvalues. Figure 14 shows the bifurcation diagram for the solution branches that are connected to the trivial solution branch with $s \leq 10$. The value of u at a generic point is plotted against the parameter s . The trivial branch has a bifurcation with \mathbb{Z}_2 symmetry at $s = \lambda_{1,1,1} = 3$ where the MI changes from 0 to 1. The trivial branch has two

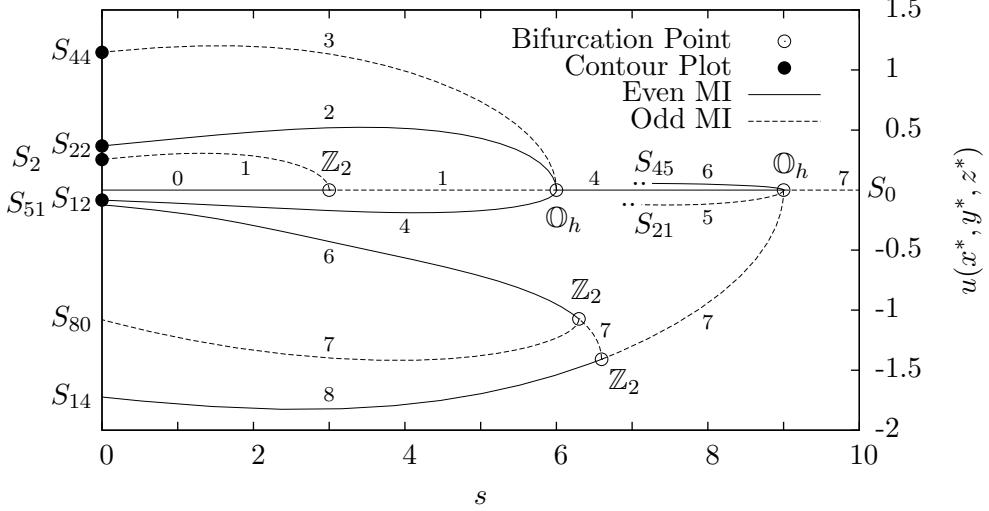


FIGURE 14. Partial bifurcation diagram for PDE (1) on the cube, showing the first three bifurcation points of the trivial branch. The two bifurcations with \mathbb{O}_h symmetry and three bifurcations with \mathbb{Z}_2 symmetry are indicated by open circles. The symmetry type of the trivial branch is S_0 , as indicated on the right. The symmetry types of the other branches are shown on the left. The small numbers indicate the Morse Index of the solution. The solution branches with symmetry type S_{21} and S_{45} are truncated to simplify the diagram. The solid dots at the left correspond to the contour maps shown in Figure 15, and contour maps for the other solution branches are shown in Figure 16.

bifurcations with \mathbb{O}_h symmetry, at $s = \lambda_{1,1,2} = 6$ and at $s = \lambda_{1,2,2} = 9$. The MI changes by 3 at each of these bifurcations, indicating a 3-dimensional critical eigenspace \tilde{E} .

At $s = 3$, the symmetry of the mother branch is $\Gamma_0 = \mathbb{O}_h \times \mathbb{Z}_2$. The critical eigenspace is $\tilde{E} = \text{span}\{\psi_{1,1,1}\}$. All of the reflections and rotations in \mathbb{O}_h act trivially on \tilde{E} . That is, $\Gamma'_0 = \langle(R_{90}, 1), (R_{120}, 1), (R_{180}, 1), (I_3, 1)\rangle = \Gamma_2 \cong \mathbb{O}_h$. Thus, the effective symmetry of the bifurcation is $\Gamma_0/\Gamma'_0 = \langle(I_3, -1)\Gamma'_0\rangle \cong \mathbb{Z}_2 = \{-1, 1\}$. The primary branch created at $s = 3$ has symmetry Γ_2 , and symmetry type $S_2 = \{\Gamma_2\}$, as shown in Figures 14 and 15.

The trivial solution undergoes a bifurcation with \mathbb{O}_h symmetry at both $s = \lambda_{1,1,2} = 6$ and $s = \lambda_{1,2,2} = 9$. The bifurcations are very similar. The critical eigenspaces for the two bifurcations are

$$\tilde{E}_6 = \text{span}\{\psi_{2,1,1}, \psi_{1,2,1}, \psi_{1,1,2}\}, \text{ and } \tilde{E}_9 = \text{span}\{\psi_{1,2,2}, \psi_{2,1,2}, \psi_{2,2,1}\},$$

respectively. The kernel of the action of $\Gamma_0 = \mathbb{O}_h \times \mathbb{Z}_2$ on \tilde{E}_6 is $\Gamma'_0 = \langle(I_3, -1)\rangle$ and the kernel of the action of Γ_0 on \tilde{E}_9 is $\Gamma'_0 = \langle(-I_3, -1)\rangle$. In both cases,

$$\Gamma_0/\Gamma'_0 = \langle(R_{90}, 1)\Gamma'_0, (R_{120}, 1)\Gamma'_0, (R_{180}, 1)\Gamma'_0\rangle \cong \mathbb{O}_h.$$

The reduced gradient $\tilde{g} : \tilde{E} \rightarrow \tilde{E}$ in both cases has the equivariance indicated in the first image of Figure 3. The Equivariant Branching Lemma (EBL) guarantees that under certain non-degeneracy conditions there is a bifurcating branch tangent to each one-dimensional fixed-point spaces in \tilde{E} . These fixed-point subspaces intersect a cube in \tilde{E} at the center of a face, the center of an edge, or a vertex of the cube. Thus, each EBL branch is made up of face, edge, or vertex solutions. The standard choice of representative in each symmetry type is $\Gamma_i \in S_i$, for $0 \leq i \leq 98$. The contour plots of the bifurcating solutions in Figure 15 show the solution in the fixed-point subspace

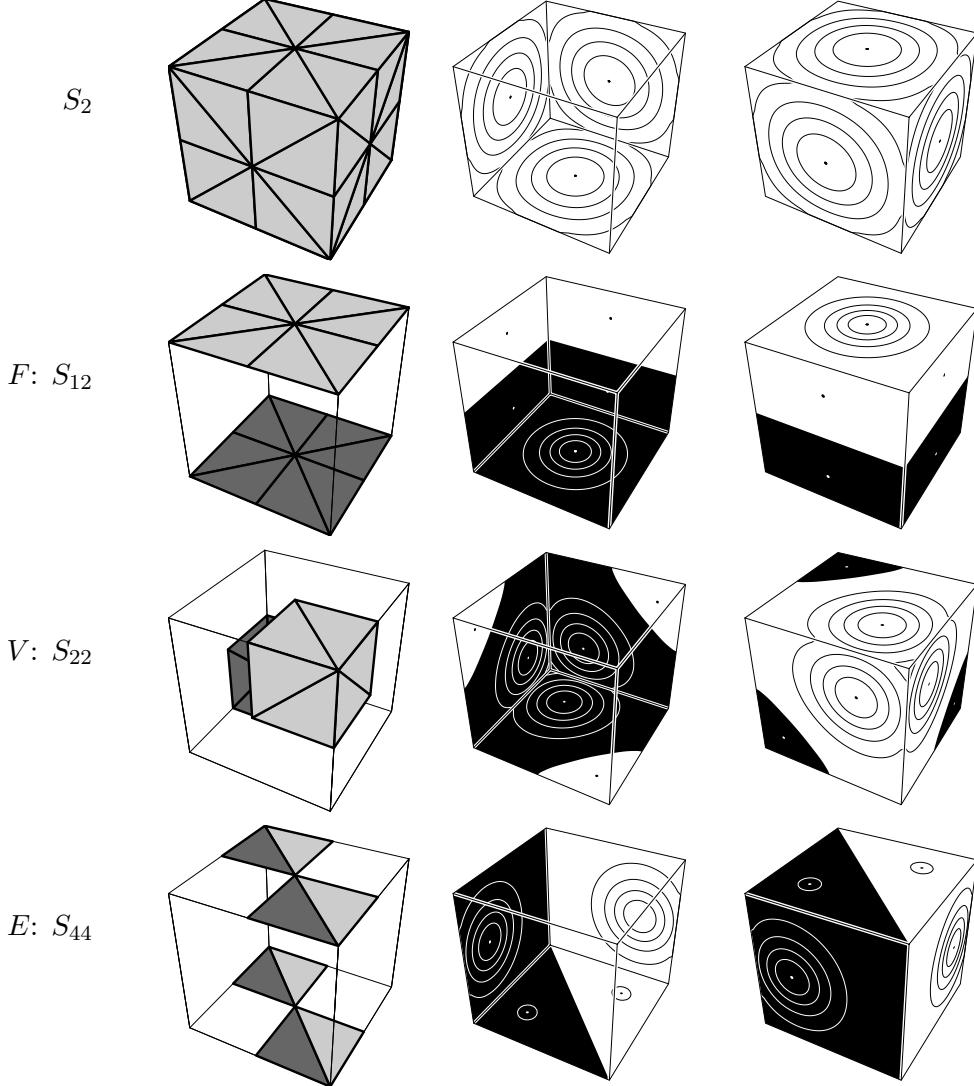


FIGURE 15. Contour plots for the solutions indicated with a solid dot in Figure 14. The contour lines for functions on the cube are equally spaced. The positive solution, with symmetry type S_2 , bifurcates at $s = 3$. Representatives of the face, vertex, and edge solutions (denoted by F , V and E) that bifurcate at $s = 6$ are shown. The vertex solution with symmetry type S_{22} is the CCN solution. The names derive for the position of the critical points of $\tilde{g} : \tilde{E} \rightarrow \tilde{E}$. The equivariance of the vector field \tilde{g} is shown in Figure 3.

indicated here:

$$\begin{aligned} \text{face } [\text{Fix}(\Gamma_{12}, \tilde{E}_6)] &= [\text{Fix}(\Gamma_{14}, \tilde{E}_9)] = \{(0, 0, a) \mid a \in \mathbb{R}\}, \\ \text{vertex } [\text{Fix}(\Gamma_{22}, \tilde{E}_6)] &= [\text{Fix}(\Gamma_{21}, \tilde{E}_9)] = \{(a, a, a) \mid a \in \mathbb{R}\}, \\ \text{edge } [\text{Fix}(\Gamma_{44}, \tilde{E}_6)] &= \{(-a, a, 0) \mid a \in \mathbb{R}\} \text{ is conjugate to} \\ &[\text{Fix}(\Gamma_{45}, \tilde{E}_9)] = \{(a, a, 0) \mid a \in \mathbb{R}\}. \end{aligned}$$

The symmetry type containing Γ_{12} , denoted S_{12} , has 3 elements, and the three conjugate face directions in $[\tilde{E}]$ are the coordinate axes. Similarly, $[\Gamma_{22}] = S_{22}$ has 4 elements, corresponding to the four diagonals through vertices of the cube centered in \mathbb{R}^3 . The edge solutions bifurcating at $s = 6$ have symmetry type $[\Gamma_{44}] = S_{44}$, which has 6 elements.

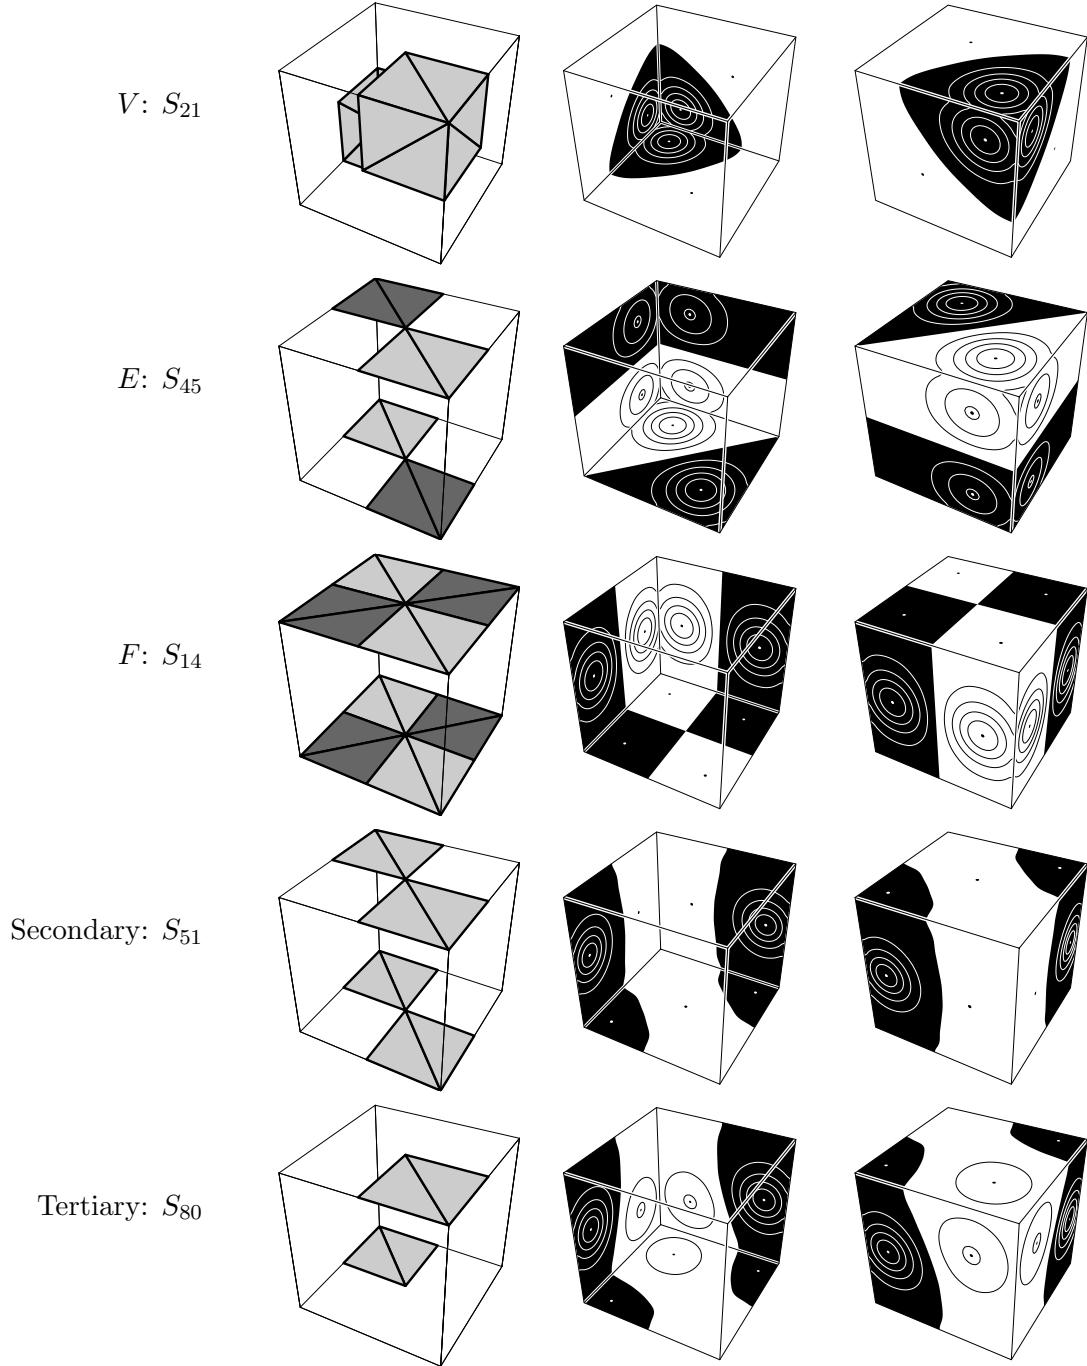


FIGURE 16. Contour plots for the vertex, edge, and face solutions that bifurcate at $s = 9$ in Figure 14, along with a daughter and granddaughter of the face solution. The contour plots are for the solutions at $s = 0$.

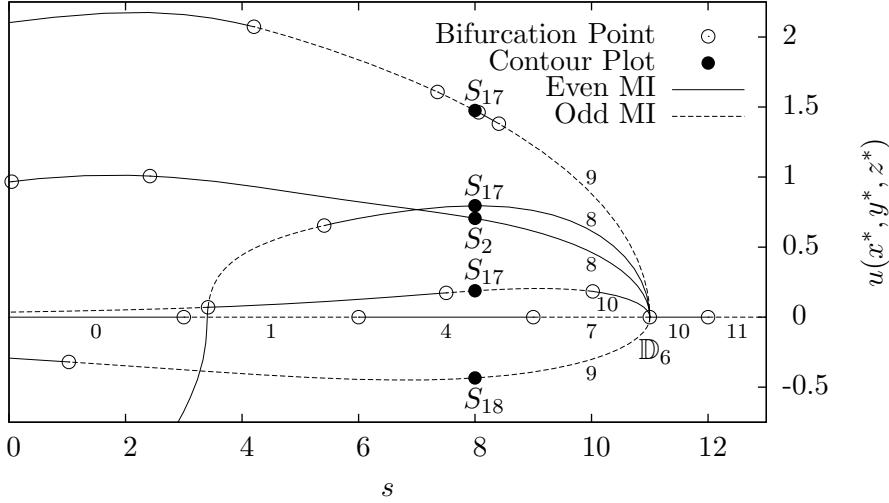


FIGURE 17. Bifurcation with \mathbb{D}_6 symmetry in PDE (1) on the cube, at $s = 11$. The symmetry of \tilde{g} on the critical eigenspace \tilde{E} is shown in Figure 18. There are 5 (conjugacy classes of) branches that bifurcate at $s = 11$. One branch in each class is shown in this bifurcation diagram, and the solid dots indicate the solutions with the contour plots in Figure 19. (The upper solid dot is very close to a bifurcation point.) There is an anomaly-breaking bifurcation at $s \approx 3.417$, where mother and daughter both have symmetry type S_{17} .

For the bifurcation at $s = 6$, the geometry in \tilde{E}_6 is mirrored in the geometry of the solutions. For example, the maximum u value for a vertex solution (type S_{22}) lies on the line from the origin to a vertex in Ω . Similarly, the maximum u value for an edge or face solution is on the line from the origin to an edge or face, respectively. The three face directions in \tilde{E} are $\psi_{2,1,1}$, $\psi_{1,2,1}$, and $\psi_{1,1,2}$, which can be thought of as “ x ”, “ y ”, and “ z ” functions. Note that the face solution in Figure 15 is approximately a multiple of $\psi_{1,1,2}$.

The geometry in \tilde{E}_9 for the bifurcation at $s = 9$ is the same. The geometry of the bifurcating solutions, shown in Figure 16, is more subtle though. The face solutions have symmetry type S_{14} , and the “ z ” eigenfunction is $\psi_{2,2,1}$. The face and edge solutions have a line where two nodal planes intersect at right angles, and these lines intersect the midpoint of a face and edge of Ω , respectively. However, the vertex solutions do not have an intersection of nodal planes. Instead, the vertex solutions have an axis of three-fold symmetry that intersects a vertex in Ω .

The face solutions (type S_{14}) that bifurcate at $s = 9$ have a bifurcation at $s \approx 6.60$, as seen in Figure 14. The secondary branch (type S_{51}) itself has a bifurcation that creates a tertiary branch with type S_{80} . Solutions from these new branches are shown in Figure 16.

5.3. A degenerate bifurcation with \mathbb{D}_6 symmetry. Figure 17 shows the bifurcation diagram of the primary branches that bifurcate from the fourth eigenvalue $s = \lambda_{3,1,1} = 11$. This is a bifurcation with \mathbb{D}_6 symmetry; the action of \mathbb{D}_6 on the critical eigenspace \tilde{E} is shown in Figure 18. Contour plots of the primary branches that bifurcate at $s = 11$ are shown in Figure 19.

The critical eigenspace is the three-dimensional space $\tilde{E} = \text{span}\{\psi_{3,1,1}, \psi_{1,3,1}, \psi_{1,1,3}\}$. The action of Γ_0 on \tilde{E} satisfies

$$\Gamma_0/\Gamma'_0 = \langle (R_{90}, 1)\Gamma'_0, (R_{120}, 1)\Gamma'_0, (I_3, -1)\Gamma'_0 \rangle \cong \mathbb{D}_6.$$

The action of Γ_0/Γ'_0 on $\tilde{E} \cong \mathbb{R}^3$ is isomorphic to the natural action of $\langle M, R_{120}, -I_3 \rangle$ on \mathbb{R}^3 , where

$$M = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

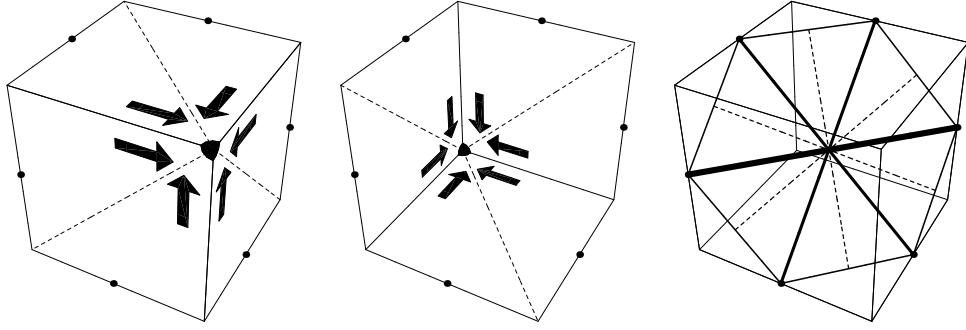


FIGURE 18. The critical eigenspace \tilde{E} at the bifurcation with \mathbb{D}_6 symmetry of the trivial solution at $s = 11$. The first two figures are similar to those of Figure 3. The arrows show the symmetry of the reduced gradient in \tilde{E} . This is a degenerate bifurcation since \tilde{E} is not an irreducible representation space. One of the diagonals of the cube in \tilde{E} is an irreducible subspace. The orthogonal subspace, which intersects the cube in a hexagon as shown in the third figure, is another irreducible subspace.

Note that M acts as a reflection. The matrix group $\langle M, R_{120}, -I_3 \rangle$ is called \mathbb{D}_{3d} in the Schönflies notation [17] for crystallographic point groups.

We now describe the symmetries of the solutions shown in Figure 19. For convenience, we will let Γ_i denote the symmetry of the solution shown with symmetry type S_i .

There is a one-dimensional fixed-point subspace of \tilde{E} for the only symmetry in $S_2 = \{\Gamma_2\}$:

$$\text{Fix}(\Gamma_2, \tilde{E}) = \text{span}\{\psi_{3,1,1} + \psi_{1,3,1} + \psi_{1,1,3}\}.$$

This space is the line through the front and back vertices shown as large dots in the first two cubes of Figure 18. The EBL guarantees that there is a solution with this symmetry; one such branch is shown in Figure 17. The other, negative, branch is not shown. Figure 19 shows a contour map of this solution with $\Gamma_2 \cong \mathbb{O}_h$ symmetry. The solution has one sign on the shaved cube, as shown, but the sign is opposite at the center of the cube. The nodal surface has cubic symmetry and is diffeomorphic to a sphere.

Figure 19 also shows one solution with symmetry $\Gamma_{18} \in S_{18}$. The one-dimensional fixed-point subspace of \tilde{E} for this symmetry is

$$\text{Fix}(\Gamma_{18}, \tilde{E}) = \text{span}\{\psi_{3,1,1} - \psi_{1,3,1}\}.$$

This fixed-point subspace is the line through the midpoints of two opposite edges, depicted as the thickest hexagon diagonal in the third cube of Figure 18. The two other diagonals of the hexagon are conjugate fixed-point subspaces. The 180° rotation about each of these diagonals is a symmetry of \tilde{E} . There is one conjugacy class of branches that bifurcates with symmetry type S_{18} at $s = 11$. There are 6 branches in this conjugacy class, one of which is shown in Figure 17.

Figure 19 shows three solutions with symmetry $\Gamma_{17} \in S_{17}$. There can be more than one conjugacy class of branches because the fixed-point subspace of the Γ_{17} action on \tilde{E} is two-dimensional:

$$\text{Fix}(\Gamma_{17}, \tilde{E}) = \text{span}\{\psi_{3,1,1} + \psi_{1,3,1}, \psi_{1,1,3}\}.$$

The intersection of this plane with a cube in \tilde{E} is indicated by dotted lines in Figure 18. This fixed-point subspace includes the one-dimensional intersection of the AIS $\mathcal{A}_{1,1,3}$ with \tilde{E} ,

$$\text{span}\{\psi_{1,1,3}\} \subseteq \text{Fix}(\Gamma_{17}, \tilde{E}),$$

so there is a bifurcating solution branch, which is approximately a multiple of $\psi_{1,1,3}$, in $\mathcal{A}_{1,1,3}$ (see Equation (11)). It is clear from the contour map that the fourth branch from the top has solutions that are in $\mathcal{A}_{1,1,3}$. Note that this branch undergoes an anomaly-breaking bifurcation at $s \approx 3.417$, with a daughter branch that has the same Γ_{17} symmetry.

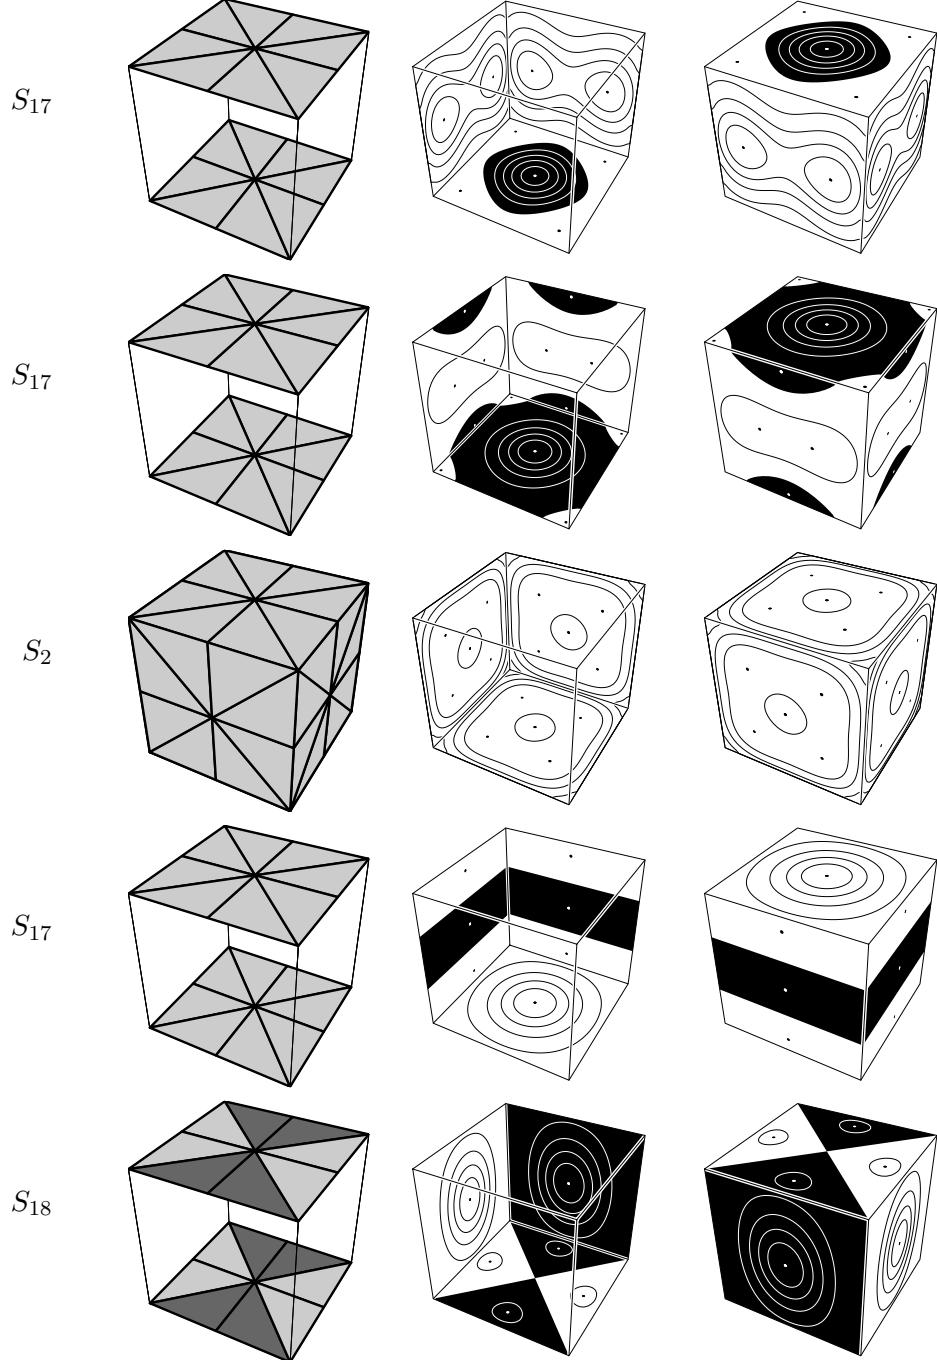


FIGURE 19. Contour plots of solutions that bifurcate at $s = 11$. The order of the solutions is the same as that in Figure 17, from top to bottom. The second solution branch intersects the fourth solution branch at an anomaly-breaking bifurcation. As described in Figure 17, both have symmetry type S_{17} . The fourth solution's contour plot shows a function which is the negative of the continuation of the second solution, due to the ink-saving heuristic that replaced u by $-u$ (see Section 3.4).

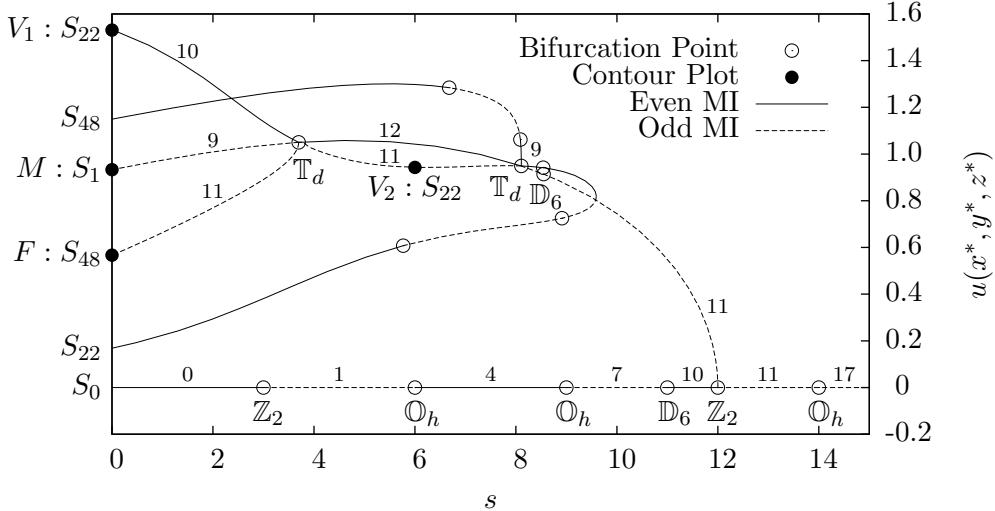


FIGURE 20. Partial bifurcation diagram for PDE (1) on the cube, showing the primary branch that bifurcates at $s = \lambda_{2,2,2} = 12$ and the two bifurcations with \mathbb{T}_d symmetry that the primary branch undergoes at $s \approx 3.697$ and $s \approx 8.107$. There is no bifurcation where the branches appear to cross but no circle is drawn. The small numerals indicate the MI of the trivial branch, the primary branch, and the solutions emanating from the bifurcation with \mathbb{T}_d symmetry on the left. Contour plots of these solutions are shown in Figure 21.

5.4. A bifurcation with \mathbb{T}_d symmetry. Our C++ program can analyze the bifurcations of nontrivial solutions, and follow all of the daughter branches of most bifurcations. For example, the primary branch that bifurcates at $s = 12$ undergoes three bifurcations in the interval $0 < s < 12$, as shown in Figure 20. Two of these bifurcations, at $s \approx 3.687$ and $s \approx 8.107$, are bifurcations with \mathbb{T}_d symmetry. The third bifurcation, at $s \approx 8.547$, is a bifurcation with \mathbb{D}_6 symmetry to be discussed later.

We focus on the bifurcation with \mathbb{T}_d symmetry at $s \approx 3.687$. The symmetry of the mother solution u^* is

$$\Gamma_1 = \langle (R_{90}, -1), (R_{120}, 1), (R_{180}, -1), (-I_3, -1) \rangle,$$

that is, $u^* \in \text{Fix}(\Gamma_1)$. For example, a rotation by 90° about the z -axis, coupled with a sign change, leaves u^* unchanged. The solution on the mother branch shown in Figure 21 looks very much like $\psi_{2,2,2}$. The critical eigenspace \tilde{E} has the ordered basis

$$(\tilde{\psi}_{2,1,1}, \tilde{\psi}_{1,2,1}, \tilde{\psi}_{1,1,2}),$$

where $\tilde{\psi}_{i,j,k}$ is a function with the same symmetry as that of $\psi_{i,j,k}$. We cannot find u^* or the critical eigenfunctions exactly with pencil and paper, but we do know the symmetry exactly, and our C++ program is able to use this information.

The action of Γ_1 on \tilde{E} satisfies $\Gamma'_1 = \langle (-I_3, -1) \rangle$ and the action of Γ_1/Γ'_1 on \tilde{E} is isomorphic to the natural action of

$$\mathbb{T}_d = \langle -R_{90}, R_{120}, -R_{180} \rangle$$

on the coordinate space $[\tilde{E}] = \mathbb{R}^3$.

The symmetry of the reduced vector field \tilde{g} on \tilde{E} for this bifurcation with \mathbb{T}_d symmetry is shown in Figure 3. The daughter solutions at this bifurcation can be classified as face solutions or vertex solutions. Each one-dimensional fixed-point subspace of the \mathbb{T}_d action on \tilde{E} is conjugate to one of these two:

$$\text{face } [\text{Fix}(\Gamma_{48}, \tilde{E})] = \{(0, 0, a) \mid a \in \mathbb{R}\}, \quad \text{vertex } [\text{Fix}(\Gamma_{22}, \tilde{E})] = \{(a, a, a) \mid a \in \mathbb{R}\}.$$

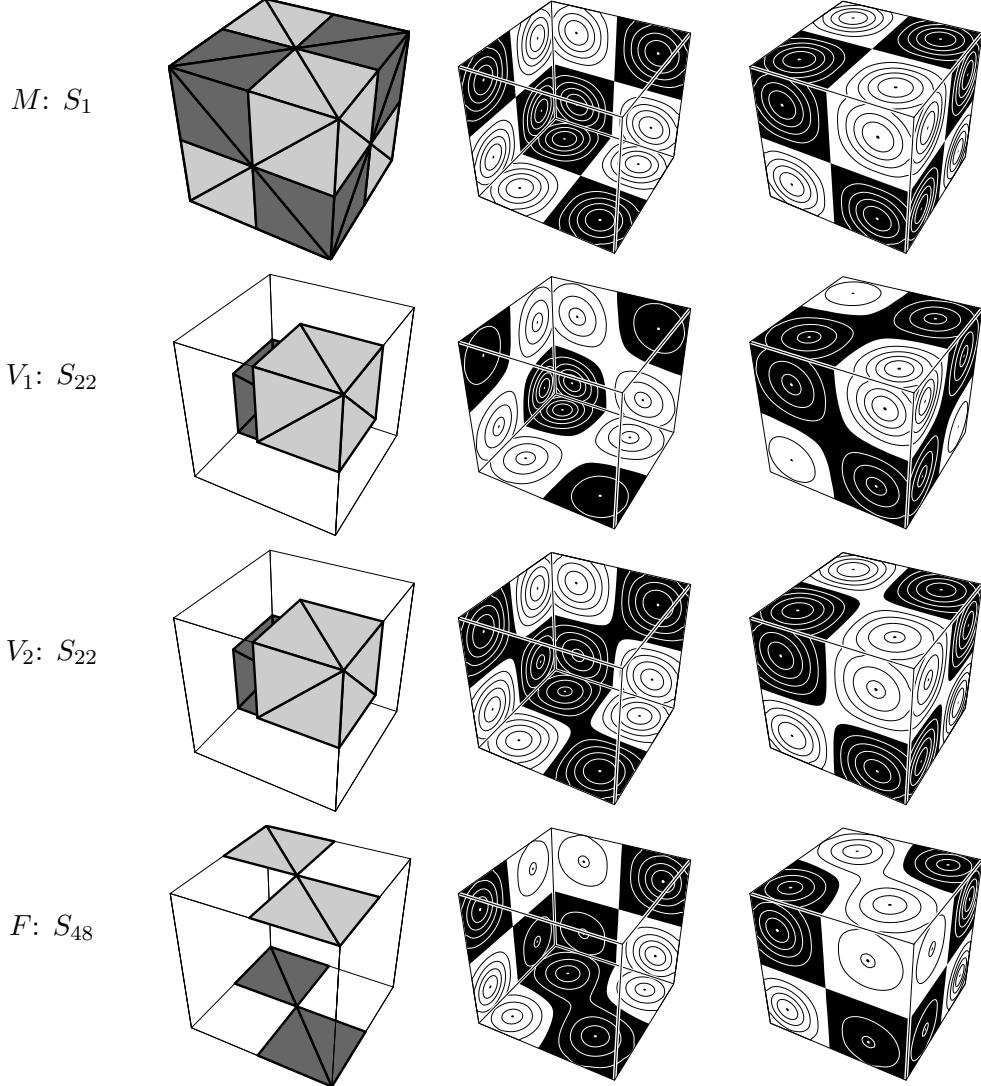


FIGURE 21. Contour plot of the mother solution M shown in Figure 20, along with two vertex solutions and a face solution born at a bifurcation with \mathbb{T}_d symmetry. The names vertex and face indicate the position in the critical eigenspace, shown in the middle image of Figure 3.

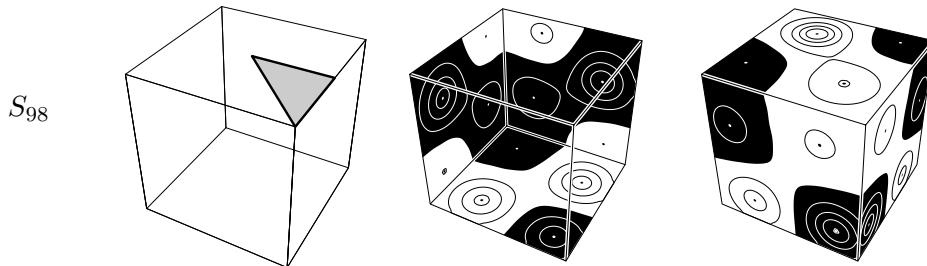


FIGURE 22. Contour plot for a solution with trivial symmetry to PDE (1) on the cube, at $s = 0$. This solution is a descendent of the mother branch with symmetry type S_1 shown in Figures 20 and 21. The solution shown is found by following the sequence of bifurcations $S_0 \xrightarrow{\mathbb{Z}_2} S_1 \xrightarrow{\mathbb{D}_6} S_{37} \xrightarrow{\mathbb{D}_4} S_{91} \xrightarrow{\mathbb{Z}_2} S_{98}$.

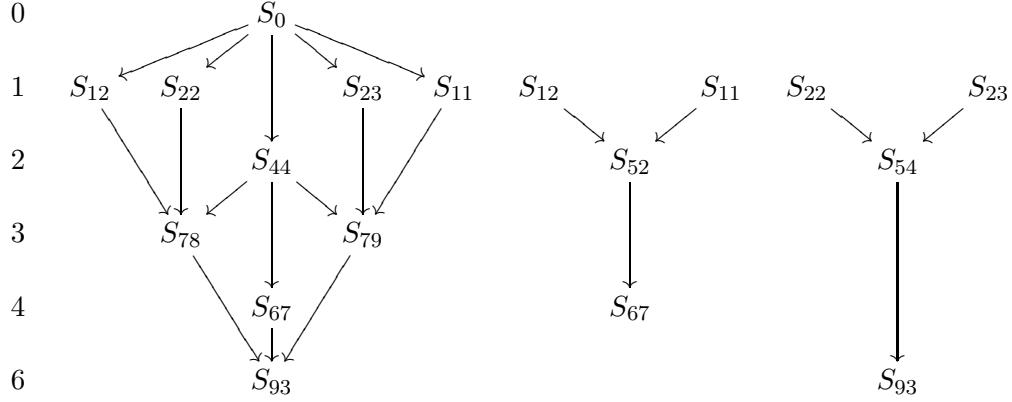


FIGURE 23. The lattice of isotropy subgroups for the six-dimensional critical eigenspace \tilde{E} of the trivial solution at $s = 14$. For clarity, we display a 3-element partition of the edge set of the Hasse diagram of the lattice. The number at the left indicates the dimension $\dim(\text{Fix}(\Gamma, \tilde{E}))$ of the fixed-point subspace for any $\Gamma \in S_i$ at that height in the diagram.

i	0	12	11	22	23	52	54	93	78	79	44	67
$ \Gamma_i \cdot u $	1	6	6	8	8	12	16	48	24	24	12	24
MI	11	12	14	14	12	15	17	16	14,15	13,16	12,13,14,15	13,14,15,17

FIGURE 24. Symmetry, multiplicity and MI at $s = s^- = 14 - \varepsilon$ of the trivial solution and the bifurcating primary solutions at $s = 14$. The solutions in a given column have symmetry type S_i . The second row shows the size of the group orbits, while the third row gives the MI of the solutions in each group orbit. The only local solution at $s = s^+ = 14 + \varepsilon$ is the trivial solution, with MI = 17. Using these MI values, we can make an index theory computation (15) to verify that the results are consistent with having obtained all solutions.

Note that I_3 , the inversion through the origin, is not in \mathbb{T}_d . In particular, two antipodal vertex solutions are not conjugate, and there is a transcritical branch of vertex solutions, as seen in Figure 20, leading to the vertex solutions V_1 and V_2 seen in Figure 21. Note that V_1 has two white regions and two black regions on the surface of the cube, whereas V_2 has one white region and one black region.

Figure 20 shows that the vertex solution V_2 is a daughter of both of the bifurcations with \mathbb{T}_d symmetry on the primary branch that bifurcates at $s = 12$. The third bifurcation on that branch, at $s \approx 8.547$, is a generic bifurcation with \mathbb{D}_6 symmetry. The MI of the mother branch changes from 11 to 9 as s decreases through that bifurcation. Unlike the degenerate bifurcation with \mathbb{D}_6 symmetry that occurs at $u = 0, s = 11$, generic bifurcations with \mathbb{D}_6 symmetry are well-known. Hence, we do not give the details of the bifurcation at $s \approx 8.547$, except to mention that one of the two branches created at this bifurcation has a grand-daughter with trivial symmetry, depicted in Figure 22.

5.5. A six-dimensional critical eigenspace. Figures 23, 24 and 25 concern the bifurcation of the trivial solution at $s = \lambda_{1,2,3} = 1^2 + 2^2 + 3^2 = 14$. The six-dimensional critical eigenspace is

$$\tilde{E} = \text{span}\{\psi_{1,2,3}, \psi_{1,3,2}, \psi_{2,1,3}, \psi_{2,3,1}, \psi_{3,1,2}, \psi_{3,2,1}\}.$$

The action of Γ_0 on \tilde{E} satisfies $\Gamma'_0 = \langle(-I_3, -1)\rangle$ and $\Gamma_0/\Gamma'_0 = \mathbb{O}_h$. The action of Γ_0/Γ'_0 on \tilde{E} is isomorphic to the natural action of

$$\langle R_{90} \oplus (-R_{90}), R_{120} \oplus R_{120}, R_{180} \oplus (-R_{180}), (-I_3) \oplus (-I_3) \rangle$$

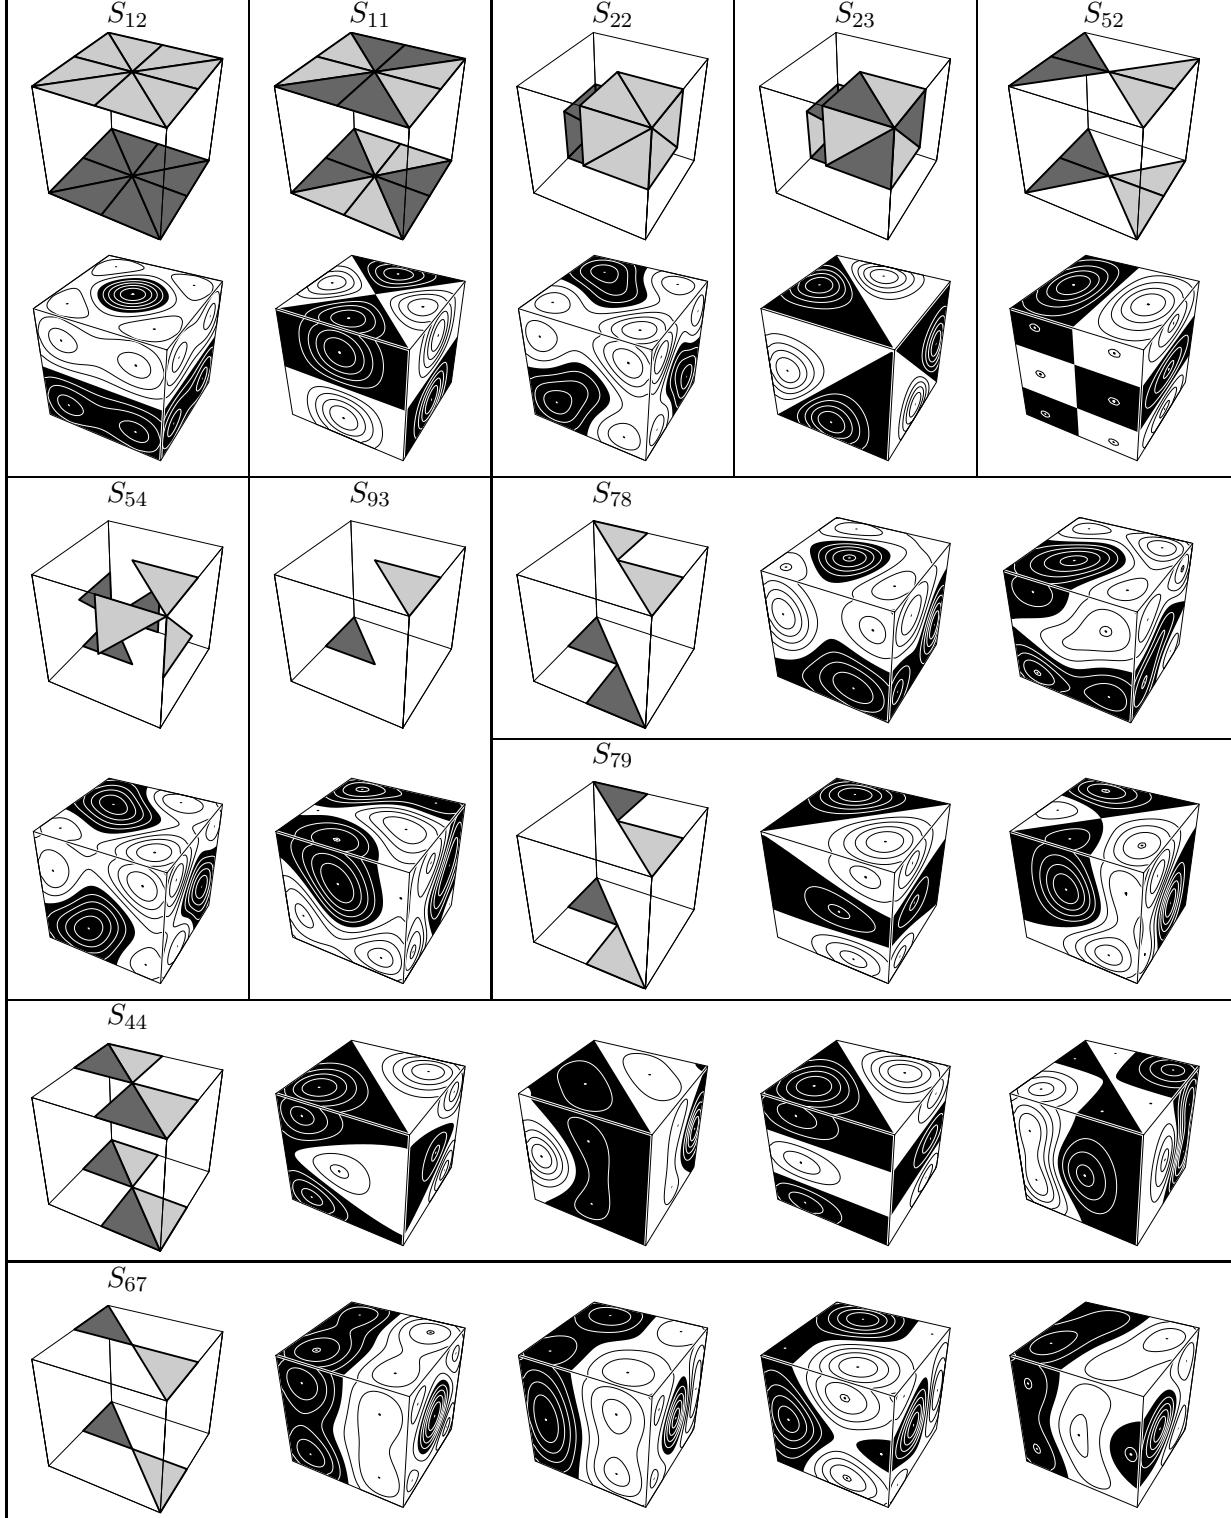


FIGURE 25. Contour plots for one solution on each of the 19 non-conjugate primary branches bifurcating at $s = 14$. The solutions are listed with increasing MI within each symmetry type. The solutions are shown at $s = 11$, except for S_{67} with MI 13 and S_{93} , which are shown at $s = 13.95$ and $s = 13.69$, respectively. We do this because these branches end at $s \approx 13.91$ and $s \approx 13.38$, respectively.

on the coordinate space $[\tilde{E}] = \mathbb{R}^6$ with respect to the ordered basis

$$(\psi_{213} + \psi_{231}, \psi_{321} + \psi_{123}, \psi_{132} + \psi_{312}, \psi_{213} - \psi_{231}, \psi_{321} - \psi_{123}, \psi_{132} - \psi_{312}).$$

Writing the action in block diagonal form, one sees that the eigenspace \tilde{E} is the direct sum of two irreducible spaces. The trivial subspace of \tilde{E} has isotropy S_0 , and S_{93} is the minimal isotropy subgroup. Thus, $[\text{Fix}(\Gamma_0, \tilde{E})] = \{0\} \subseteq \mathbb{R}^6$ and $[\text{Fix}(\Gamma_{93}, \tilde{E})] = \mathbb{R}^6$. The remaining symmetries in Figures 23 and 25 in our chosen coordinate space satisfy the following:

$$\begin{aligned} [\text{Fix}(\Gamma_{12}, \tilde{E})] &= \{(0, 0, a, 0, 0, 0) \mid a \in \mathbb{R}\}, & [\text{Fix}(\Gamma_{11}, \tilde{E})] &= \{(0, 0, 0, 0, 0, a) \mid a \in \mathbb{R}\}, \\ [\text{Fix}(\Gamma_{22}, \tilde{E})] &= \{(a, a, a, 0, 0, 0) \mid a \in \mathbb{R}\}, & [\text{Fix}(\Gamma_{23}, \tilde{E})] &= \{(0, 0, 0, a, a, a) \mid a \in \mathbb{R}\}, \\ [\text{Fix}(\Gamma_{52}, \tilde{E})] &= \{(0, a, 0, 0, b, 0) \mid a, b \in \mathbb{R}\}, & [\text{Fix}(\Gamma_{54}, \tilde{E})] &= \{(a, a, a, b, b, b) \mid a, b \in \mathbb{R}\}, \\ [\text{Fix}(\Gamma_{44}, \tilde{E})] &= \{(a, -a, 0, b, b, 0) \mid a, b \in \mathbb{R}\}, & [\text{Fix}(\Gamma_{78}, \tilde{E})] &= \{(a, -a, b, c, -c, 0) \mid a, b, c \in \mathbb{R}\}, \\ [\text{Fix}(\Gamma_{67}, \tilde{E})] &= \{(a, b, 0, c, d, 0) \mid a, b, c, d \in \mathbb{R}\}, & [\text{Fix}(\Gamma_{79}, \tilde{E})] &= \{(a, a, 0, b, b, c) \mid a, b, c \in \mathbb{R}\}. \end{aligned}$$

Figure 23 describes the lattice of isotropy subgroups of \tilde{E} . Each arrow $S_i \rightarrow S_j$ indicates that some isotropy subgroup in S_j is a subgroup of some isotropy subgroup in S_i . The arrows generate a partial ordering of the symmetry types. Note that the lattice of isotropy subgroups is different from the bifurcation digraph, as explained in [13].

To simplify the visual representation, the lattice of symmetry types for the action of \mathbb{O}_h on two irreducible spaces whose direct sum is \tilde{E} are shown on the top row of Figure 23. Note that the middle column is the same in each of the top row sub-lattices. As a result of the presence of S_{44} in both sub-lattices, $\dim(\tilde{E} \cap \text{Fix}(\Gamma_{44})) = 2$, whereas $\dim(\tilde{E} \cap \text{Fix}(\Gamma_i)) = 1$ for $i \in \{11, 12, 22, 23\}$. Within these one-dimensional spaces there is a pitchfork bifurcation to an EBL branch, but the bifurcation to solutions with symmetry type S_{44} is more complicated.

It is remarkable that there is at least one solution branch bifurcating at $s = 14$ with each of the symmetry types shown in Figure 23. There is even a solution with symmetry type S_{93} , the lowest symmetry present in \tilde{E} . The conjugacy class of this branch has a total of 48 branches. Figure 25 shows one solution in each of the nonconjugate primary branches that bifurcate at the multiplicity six eigenvalue $s = 14$. Since each of the solutions in this figure is odd about the center of the cube, that is $u(x, y, z) = -u(\pi - x, \pi - y, \pi - z)$, we only show the front view of the contour plot.

The solution in Figure 25 with symmetry type S_{52} strongly resembles the eigenfunction $\psi_{1,2,3}(x, y, z) = \sin(x) \sin(2y) \sin(3z)$. An analysis of the “hidden symmetries” in this problem [10] would explain why the solution with symmetry type S_{52} bifurcates, but it would not explain all of the solutions in Figure 25. In the space of triply periodic functions on \mathbb{R}^3 , there is a 12-dimensional irreducible space spanned by rotations of $\psi_{1,2,3}$ and the similar functions with cosines in place of sines.

Let X^- and X^+ be the set of solutions for $s = s^- = 14 - \varepsilon$ and $s = s^+ = 14 + \varepsilon$, respectively, that are on branches bifurcating from $(0, 14) \in H \times \mathbb{R}$, together with solutions on the mother branch, for a sufficiently small positive ε . The set X^- contains 345 solutions falling into 20 group orbits with non-trivial representatives shown in Figure 25. Figure 24 shows multiplicity and MI information for X^- . Since all the bifurcating branches curve to the left, X^+ contains only the trivial solution with $\text{MI}(0, s^+) = 17$. Thus, one can verify that the Poincaré-Hopf Index Theorem of [2] is satisfied since

$$(15) \quad \sum_{u \in X^-} (-1)^{\text{MI}(u, s^-)} = \sum_i \sum_{u \in X^- / \Gamma_i} |\Gamma_i \cdot u| (-1)^{\text{MI}(u, s^-)} = -1 = (-1)^{17} = \sum_{u \in X^+} (-1)^{\text{MI}(u, s^+)}.$$

This is consistent with our belief that the list of solutions in Figure 24 is comprehensive.

6. CONCLUSION

In this article we have extended the methods from [12] and [13]. In the first paper, the symmetry group was relatively small and a medium-sized grid was used to produce a reasonable portion of the

bifurcation diagram and a selection of contour plots for a two-dimensional semilinear elliptic PDE. In the second article, we completely automated the symmetry analysis for investigating the rich symmetries of solutions to partial difference equations (PdE) for many interesting low-order graphs. In the current article, we have shown how to extend these ideas to a three-dimensional problem with a large symmetry group, namely the cube. The large grid and many calculations required the use of a parallel programming environment. We developed and employed our own library, MPQueue, to implement our branch following and branch switching algorithms using self-submitting parallel job queues and MPI. The new results we have presented here use the symmetry analysis from a low-order graph with the same symmetry group to generate the corresponding symmetry information for functions discretized over the large-sized grid used in the PDE code. We could not have done this without GAP; for the cube there are 99 symmetry types (and 323 symmetries or isotropy subgroups) with 482 arrows between symmetry types. This symmetry information is essential to the numerical results in several key ways. It allows for the efficient construction of block-diagonal Hessians, reducing the number of costly integrations required at each Newton step in the presence of symmetry. It allows us to search for only a single representative of each novel solution type, rather than wasting computations on finding many equivalent copies. By reducing the dimension of a search space, symmetry information increases our chance of finding all expected solutions of a given symmetry type at each new bifurcation. The entire suite of programs and new methods for efficiently implementing our algorithms has allowed us to observe interesting bifurcation symmetries for PDE that we have not previously seen published. Our procedure demonstrates a robustness for handling degenerate bifurcations, AIS, and high-dimensional/reducible critical eigenspaces. The new contour plots required a number of ideas for efficiently and effectively conveying the necessary information graphically. The size of the problem makes it impossible to present a visual representation of the bifurcation digraph on a single page. We have constructed a companion website for navigating the digraph, and give examples here to aid the reader in understanding the digraph and how to use it to interpret our numerical results.

REFERENCES

- [1] Antonio Ambrosetti and Paul H. Rabinowitz. Dual variational methods in critical point theory and applications. *J. Functional Analysis*, 14:349–381, 1973.
- [2] Vladimir I. Arnol'd. *Ordinary differential equations*. Springer Textbook. Springer-Verlag, Berlin, 1992. Translated from the third Russian edition by Roger Cooke.
- [3] C. J. Budd, A. R. Humphries, and A. J. Wathen. The finite element approximation of semilinear elliptic partial differential equations with critical exponents in the cube. *SIAM J. Sci. Comput.*, 20(5):1875–1904, 1999.
- [4] Alfonso Castro, Jorge Cossio, and John M. Neuberger. A sign-changing solution for a superlinear Dirichlet problem. *Rocky Mountain J. Math.*, 27(4):1041–1053, 1997.
- [5] Goong Chen and Jianxin Zhou. *Boundary element methods with applications to nonlinear problems*, volume 7 of *Atlantis Studies in Mathematics for Engineering and Science*. Atlantis Press, Paris, second edition, 2010.
- [6] MPI Forum. MPI: A message-passing interface standard. version 1.x. <http://www.mpi-forum.org>.
- [7] Victor A. Galaktionov and Sergey R. Svirshchevskii. *Exact solutions and invariant subspaces of nonlinear partial differential equations in mechanics and physics*. Chapman & Hall/CRC Applied Mathematics and Nonlinear Science Series. Chapman & Hall/CRC, Boca Raton, FL, 2007.
- [8] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.4.9*, 2006.
- [9] Martin Golubitsky, Ian Stewart, and David G. Schaeffer. *Singularities and groups in bifurcation theory. Vol. II*, volume 69 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 1988.
- [10] M. Gabriela M. Gomes, Isabel S. Labouriau, and Eliana M. Pinho. Spatial hidden symmetries in pattern formation. In *Pattern formation in continuous and coupled systems (Minneapolis, MN, 1998)*, volume 115 of *IMA Vol. Math. Appl.*, pages 83–99. Springer, New York, 1999.
- [11] John M. Neuberger, Nándor Sieben, and James W. Swift. Computing eigenfunctions on the Koch snowflake: a new grid and symmetry. *J. Comput. Appl. Math.*, 191(1):126–142, 2006.
- [12] John M. Neuberger, Nándor Sieben, and James W. Swift. Symmetry and automated branch following for a semilinear elliptic PDE on a fractal region. *SIAM J. Appl. Dyn. Syst.*, 5(3):476–507 (electronic), 2006.
- [13] John M. Neuberger, Nándor Sieben, and James W. Swift. Automated bifurcation analysis for nonlinear elliptic partial difference equations on graphs. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 19(8):2531–2556, 2009.

- [14] John M. Neuberger, Nándor Sieben, and James W. Swift. Companion web site, 2012. <http://www.nau.edu/Jim.Swift/cubePDE>.
- [15] John M. Neuberger, Nándor Sieben, and James W. Swift. An MPI implementation of a self-submitting parallel job queue. *International Journal of Parallel Programming*, 40(4):443–464, 2012.
- [16] John M. Neuberger and James W. Swift. Newton’s method and Morse index for semilinear elliptic PDEs. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 11(3):801–820, 2001.
- [17] Michael Tinkham. *Group theory and quantum mechanics*. McGraw-Hill Book Co., New York, 1964.
- [18] Zhi-Qiang Wang and Jianxin Zhou. A local minimax-Newton method for finding multiple saddle points with symmetries. *SIAM J. Numer. Anal.*, 42(4):1745–1759 (electronic), 2004.
- [19] Zhi-Qiang Wang and Jianxin Zhou. An efficient and stable method for computing multiple saddle points with symmetries. *SIAM J. Numer. Anal.*, 43(2):891–907 (electronic), 2005.
- [20] Jianxin Zhou. Saddle critical point analysis and computation. *Nonlinear Anal.*, 63(5-7):1000–1009, 2005.

E-mail address: John.Neuberger@nau.edu, Nandor.Sieben@nau.edu, Jim.Swift@nau.edu

DEPARTMENT OF MATHEMATICS AND STATISTICS, NORTHERN ARIZONA UNIVERSITY PO BOX 5717, FLAGSTAFF, AZ 86011-5717, USA